

2 Staring Emmy Straight in the Eye—And Doing My Best Not to Flinch

Douglas Hofstadter

Editor's Introduction

Doug Hofstadter describes himself modestly as professor of cognitive science at the College of Arts and Sciences at Indiana University Bloomington (where he also directs the Center for Research on Concepts and Cognition) and someone who has had a life-long love for and involvement in music. Beyond this, of course, Hofstadter is an undisputed major figure in cognitive science and especially in the study of creativity. His Pulitzer Prize-winning book Gödel, Escher, Bach and numerous more recent books cover vast areas of interdisciplinary studies of the visual arts, artificial intelligence, language, music, and mathematics.

Hofstadter here offers a number of critical insights into what he believes are the implications and challenges of the Experiments in Musical Intelligence program. He shares with us his angst over the program's apparent ability to fool people with its virtual music and to often provide convincing musical experiences which he heretofore believed only human-composed music could inspire.

As Doug states, he and I have had many discussions about Experiments in Musical Intelligence. His article here demonstrates a keen understanding of the basic fundamentals of the program. However, Doug's interpretation of the program does not include a discussion of earmarks, transformation, and so on due either to their recent appearance in the program or the unnecessary complication their inclusion would create. He nonetheless gives a compelling description of how Experiments in Musical Intelligence works.

For the record, Doug refers to Experiments in Musical Intelligence using the familiar name "Emmy" in his prose and "EMI" or "E.M.I." in his poetry. I have preserved the latter acronym for purposes of rhyme but wish to make clear that there is no relationship between the EMI or E.M.I. acronyms and the Thorne EMI Corporation.

Good artists borrow; great artists steal.
—Douglas Hofstadter¹

How Young I Was, and How Naive

I am not now, nor have I ever been, a card-carrying futurologist. I make no claims to be able to peer into the murky crystal ball and make out what lies far ahead. But one time, back in 1977, I did go a little bit out on a futurologist's limb. At the end of

Chapter 19 (“Artificial Intelligence: Prospects”) of my book *Gödel, Escher, Bach* (1979), I had a section called “Ten Questions and Speculations,” and in it I stuck my neck out, venturing a few predictions about how things would go in the development of AI. Though it is a little embarrassing to me now, let me nonetheless quote a few lines from that section here:

Question: Will there be chess programs that can beat anyone?

Speculation: No. There may be programs which can beat anyone at chess, but they will not be exclusively chess players. They will be programs of *general* intelligence, and they will be just as temperamental as people. “Do you want to play chess?” “No, I’m bored with chess. Let’s talk about poetry.” That may be the kind of dialogue you could have with a program that could beat everyone . . .

We all know today how very wrong that speculation was. What was it that so misled the author of *Gödel, Escher, Bach* back then?

Well, when I wrote those words, I was drawing some of my ideas from a fascinating article that I had read by my soon-to-be colleague at Indiana University, the psychologist and chess master Eliot Hearst (formerly vice president of the U.S. Chess Federation [USCF], member of the U.S. Chess Olympics team, and once a frequent playing partner of Bobby Fischer). In his article (1977), Hearst (who clearly knew infinitely more about chess than I ever could hope to) eloquently expressed the conviction that deep chess-playing ability depends in an intimate manner on such cognitive skills as the ability to sort the wheat from the chaff in an intuitive flash, the ability to make subtle analogies, and the ability to recall memories associatively. All of these elusive abilities seemed to lie so close to the core of human nature itself that I jumped to the conclusion that profoundly insightful chess-playing draws intrinsically on central facets of the human condition, and that mere brute-force searching of the rapidly branching look-ahead tree, no matter how fast, broad, or deep, would not be able to circumvent or shortcut that fact.

I didn’t realize—and perhaps no one did at the time—that the USCF rankings of the best computer chess programs (all of which used brute-force search algorithms) were pretty much creeping up linearly with time, so that a simple-minded linear extrapolation on a plot of chess prowess vs. time would, even back then, have suggested that computers would take over from humans somewhere around the year 2000. The first time I actually saw such a graph was in an article in *Scientific American* in the mid-1990s (written by the creators of Deep Blue, by the way), and I vividly remember thinking to myself, when I looked at it, “Uh-oh! The handwriting is on the wall!” And so it was.

Chess Tumbles to Computational Power ...

We now know that world-class chess-playing ability can indeed be achieved by brute-force techniques—techniques that in no way attempt to replicate or emulate what goes on in the head of a chess grandmaster. Analogy-making is not needed, nor is associative memory, nor are intuitive flashes that sort wheat from chaff—just a tremendously wide and deep search, carried out by superfast, chess-specialized hardware using ungodly amounts of stored knowledge. And thus, thanks to the remarkable achievements of the past decade, one can no longer look at a subtle, elegant, and stunning midgame chess move and say with confidence, “Only a genius could have spotted that move!” because the move could just as well have emanated from a mindless, lightning-fast full-width search as from the silent machinations of an insightful human mind.

I cannot say what goes on in the brain of a Bobby Fischer or a Garry Kasparov when they play championship-level chess. I have no idea whether their phenomenal chess-playing ability draws in some subtle way on their entire human existence, on their prior struggles with life and death, on their striving for personal identity, on their coping with dashed romances, on their hopes and fears in domains apparently remote from chess—or, contrariwise, whether their chess-playing skill is in some sense totally isolated from the rest of their minds, fully contained in some little localized region of their brains that, at least in principle, could be neatly excised by a neurosurgeon, leaving the rest of their brains fully intact so that they could go on living normal lives while the little module, safely preserved and nourished in a vat, happily kept on playing world-level chess.

Eliot Hearst’s article had led me to believe that the image of an isolated chess-playing module is wrong, and that, to the contrary, great chess-playing skill is of necessity deeply intertwined with all that being human is about. But as Deep Blue has taught us, that certainly need not be the case. Topnotch chess-playing does not necessarily depend on the full mental complexities that come from living life, facing death, and all those messy things that we experience. Topnotch chess playing *can* come from a pure chess engine, full stop. As for topnotch *human* chess-playing ability, one might still plausibly believe that it is necessarily tightly integrated with the rest of the brain and with the whole kit and caboodle of being human—but ever since Deep Blue’s appearance on the scene, there is reason to doubt that romantic vision. Perhaps it is the case, but perhaps not.

I, in any case, have had to eat humble pie with respect to my 1977 speculation. But, I must say, having to swallow my words about chess doesn’t upset me all that much, since, aside from writing that one speculation, I personally have never had any emo-

tional stake in the notion that chess skill lies very near the pinnacle of that which is most truly human, and so I'm not crushed that my speculation was refuted. And even though people say that the game of Go is far less computer-tractable than chess is, I don't think I'd care to rewrite my speculation substituting Go for chess. I'll just admit my mistake.

So . . . chess-playing fell to computers? I don't feel particularly threatened or upset; after all, sheer computation had decades earlier fallen to computers as well. So a computer had outdone Daniel Shanks in the calculation of digits of π —did it matter? Did that achievement in any way lower human dignity? Of course not! It simply taught us that calculation is more mechanical than we had realized. Likewise, Deep Blue taught us that chess is more mechanical than we had realized. These lessons serve as interesting pieces of information about various domains of expertise, but to my mind they hardly seem to threaten the notion, which I then cherished and which I still cherish, that human intelligence is extraordinarily profound and mysterious.

It is not, I hasten to add, that I am a mystic who thinks that intelligence intrinsically resists implantation in physical entities. To the contrary, I look upon brains themselves as very complex machines, and, unlike John Searle and Roger Penrose, I have always maintained that the precise nature of the physicochemical substrate of thinking and consciousness is irrelevant. I can imagine silicon-based thought as easily as I can imagine carbon-based thought; I can imagine ideas and meanings and emotions and a first-person awareness of the world (an “inner light,” a “ghost in the machine”) emerging from electronic circuitry as easily as from proteins and nucleic acids. I simply have always run on faith that when “genuine artificial intelligence” (sorry for the oxymoron) finally arises, it will do so precisely because the same degree of complexity and the same overall kind of abstract mental architecture will have come to exist in a new kind of hardware. What I do *not* expect, however, is that full human intelligence will emerge from something far simpler, architecturally speaking, than a human brain.

. . . and so, Is Musical Beauty Next in Line?

My “Ten Questions and Speculations” section in *GEB* was an attempt to articulate just these kinds of pieces of faith, and at the time I wrote it, I was particularly proud of another one of them, which I now reproduce here in full:

Question: Will a computer program ever write beautiful music?

Speculation: Yes, but not soon. Music is a language of emotions, and until programs have emotions as complex as ours, there is no way a program will write anything beautiful. There

can be “forgeries”—shallow imitations of the syntax of earlier music—but despite what one might think at first, there is much more to musical expression than can be captured in syntactical rules. There will be no new kinds of beauty turned up for a long time by computer music-composing programs. Let me carry this thought a little further. To think—and I have heard this suggested—that we might soon be able to command a preprogrammed mass-produced mail-order twenty-dollar desk-model “music box” to bring forth from its sterile [sic!] circuitry pieces which Chopin or Bach might have written had they lived longer is a grotesque and shameful misestimation of the depth of the human spirit. A “program” which could produce music as they did would have to wander around the world on its own, fighting its way through the maze of life and feeling every moment of it. It would have to understand the joy and loneliness of a chilly night wind, the longing for a cherished hand, the inaccessibility of a distant town, the heartbreak and regeneration after a human death. It would have to have known resignation and world-weariness, grief and despair, determination and victory, piety and awe. In it would have had to commingle such opposites as hope and fear, anguish and jubilation, serenity and suspense. Part and parcel of it would have to be a sense of grace, humor, rhythm, a sense of the unexpected—and of course an exquisite awareness of the magic of fresh creation. Therein, and therein only, lie the sources of meaning in music.

In recent years, when lecturing about Dave Cope’s work, I have read this paragraph aloud so many times that I practically know it by heart. And what do I make of it now? Well, I am not quite sure. I have been grappling for several years now with these issues, and still there is no clear resolution. That, perhaps, is why I have been so fascinated by Cope’s Emmy and the issues raised thereby. Let me explain.

In the spring of 1995, I was conducting a cognitive science seminar at Indiana University called “AI: Hype versus Hope,” whose purpose was for me and my students, working together, to try to sort the wheat from the chaff in this field so rife with brazen claims of human-level performance in one domain or another, most of which I knew were groundless, or nearly so. I was willing to concede, however, that even in a hopelessly hyped project, there might somewhere reside a nugget of value, and it was my idea that we would uncover those nuggets while at the same time chucking out the overblown claims. We discussed computer driving of cars, speech recognition, story understanding, machine translation, face recognition, and many other topics. One topic that particularly interested me was music, because I was convinced, a priori, that claims I’d heard here and there about high-quality music emanating from computers were hugely exaggerated, and I wanted to confirm this hunch. And so when a student in the seminar told me she had run across a book called *Computers and Musical Style* in the music library and wondered if she could present it to the seminar, I enthusiastically encouraged her to do so.

A couple of days later in class, this student described to us the ideas behind the program—Emmy, to be specific—but I found myself not terribly interested. It sounded like Emmy was dealing only with the surface level of music—with patterns,

not with the deep emotional substrate—and I was pretty sure that little of interest could come of such an architecture. Then she said she could play for us some of Emmy’s compositions on the piano in my research center, so I said “Fine!” We went in and listened as she played, and my skeptical ears were somewhat jolted. Although the two pieces she played—very short Mozart-aping and Brahms-aping pieces—sounded amateurish and flawed, they were by no means totally incoherent or absurd. I wondered how in the world they could have come out of this architecture, and so I asked if I could borrow the book for a day or two. She said yes, and I took it home and plunged into it with great interest.

I noticed in its pages an Emmy mazurka supposedly in the Chopin style, and this really drew my attention because, having revered Chopin my whole life long, I felt certain that no one could pull the wool over my eyes in this department. Moreover, I knew all fifty or sixty of the Chopin mazurkas very well, having played them dozens of times on the piano and heard them even more often on recordings. So I went straight to my own piano and sight-read through the Emmy mazurka—once, twice, three times, and more—each time with mounting confusion and surprise. Though I felt there were a few little glitches here and there, I was impressed, for the piece seemed to *express* something. If I had been told it had been written by a human, I would have had no doubts about its expressiveness. I don’t know that I would have accepted the claim that it was a newly uncovered mazurka by Chopin himself, but I would easily have believed it was by a graduate student in music who loved Chopin. It was slightly nostalgic, had a bit of Polish feeling in it, and it did not seem in any way plagiarized. It was *new*, it was unmistakably *Chopin-like* in spirit, and it was *not emotionally empty*. I was truly shaken. How could emotional music be coming out of a program that had never heard a note, never lived a moment of life, never had any emotions whatsoever?

The more I grappled with this, the more disturbed I became—but also fascinated. There was a highly counterintuitive paradox here, something that obviously had caught me enormously off guard, and it was not my style to merely deny it and denounce Emmy as “trivial” or “nonmusical.” To do so would have been cowardly and dishonest. I was going to face this paradox straight on, and it seemed to me that the best thing to do was to look the monster right in the face. And thus I picked up my telephone and phoned the program’s inventor, David Cope, in Santa Cruz. I reached him with ease, and as he was very friendly and open, I asked him about aspects of Emmy’s architecture that I had not been able to glean from his book. After a lengthy and very informative conversation, we made a point of agreeing to get together next time I was in California. In the meantime, I continued to grapple with this strange program that was threatening to upset the apple cart that held many of

my oldest and most deeply cherished beliefs about the sacredness of music, about music being the ultimate inner sanctum of the human spirit, the last thing that would tumble in AI's headlong rush toward thought, insight, and creativity.

The Proof of the Pudding Is in the Eating

Of all the projects examined in my “Hype versus Hope” seminar, Emmy was the only one that made me reconsider deeply held beliefs. I have to confess, though, that had I only read about its architecture and not heard any of its output, I would have paid little or no attention to it. Although Cope has put in far more work on Emmy than most AI researchers ever do on any one project (he was worked on it for nearly 20 years now, and the program consists of some 20,000 lines of Lisp code that runs on his trusty Macintosh), the basic ideas in the design of Emmy simply did not sound radically new to me, or even all that promising. What made all the difference in the world for me was *carefully listening to Emmy's compositions*.

I don't think one can possibly judge Emmy without hearing some of “her” pieces (Dave usually says “her,” and, for fun, I sometimes go along with the anthropomorphism). Some people will approach them open-mindedly, while others—often musicians—will come to Emmy's pieces with a strong preconceived idea that they will be weak or blatantly derivative, and so, however the pieces actually sound, such people will wind up putting them down, even pooh-poohing them, safe in their knowledge that they were generated by a computer. For that reason, I think it best that one first hear a few of Emmy's pieces without knowing their provenance—perhaps without even having ever heard of Emmy. I don't like dishonesty, but perhaps it is best to misinform people about what they are about to hear, in order that they not listen with a preclosed mind.

Lecturing on Emmy in Many Different Venues

It was not too long after my first exposure to Emmy that I decided that I had to organize my many complex reactions to this strange project in a coherent fashion, and that meant preparing a well-rounded lecture on it all. I pulled together a set of thoughts, made a bunch of transparencies, and was lucky enough to find several venues where I could give this lecture. My set of transparencies evolved in many ways as these lectures took place, which was good, but one strange thing I soon discovered was that almost no one in my various audiences shared my profound sense of bewilderment or alarm. Hardly anyone seemed upset at Cope's coup in the modeling of

artistic creativity; hardly anyone seemed threatened or worried at all. I felt kinship with but a few souls in the world who also were bewildered by similar triumphs. One of them was none other than Garry Kasparov, who had said, a year before being trounced by Deep Blue:

To some extent, this match is a defense of the whole human race. Computers play such a huge role in society. They are everywhere. But there is a frontier that they must not cross. They must not cross into the area of human creativity. It would threaten the existence of human control in such areas as art, literature, and music. (Kasparov 1996)

On one level, Kasparov's words sounded ridiculous to me. Saying computers "must not cross into . . . human creativity" seemed hopelessly naive, almost like saying, "We must not let them do certain things, because they'll beat our pants off if we do, and won't that be dreadful!" And Kasparov's last sentence, even sillier, raises the specter of computers trying to wrest control away from human beings, as if on the surface of our planet there were already raging some terrible battle between alien species for control of culture. Such a weird scenario may possibly come to be in the next few decades or next few centuries—who can say for sure?—but certainly it is not happening already. Today we control computers, and that is beyond doubt or dispute.

And yet . . . and yet . . . something of Kasparov's worried tone resonated with me. It was as if he had felt, and I now felt, something about the profundity of the human mind's sublimity being taken away, being robbed, by the facile victories of programs that seemed totally out of touch with the essence of the domains in which they were operating so well. It seemed somehow humiliating, even nightmarish, to me.

But no matter how I tried, I could not get my own sense of confusion and worry across to my audience. One thing I learned fairly soon was that few people have a visceral feeling about the centrality and depth of music. Indeed, I discovered that there is a rough trichotomy of people. There are some who, like me, feel that music is the most powerful drug in the world, and that it reaches in and touches one's innermost core like almost nothing else—more powerfully than art, than literature, than cinema, and so on. But such people are few and far between. A much more common attitude is, "Sure I like music, but it doesn't touch me at my very core. It's just fun to listen to, to dance to, and so forth." And then another attitude that came up surprisingly often in question-and-answer sessions after my lectures was this: "I'm kind of tone-deaf, and music's okay but I can take it or leave it, so I don't really relate to your deep love of music, but . . ."

I soon realized that I was probably not going to reach the third group no matter what I said, and wondered if the "music enthusiasts" of the middle group were also

beyond reach. But to my greater chagrin, even most people in the *first* group often couldn't relate to my worry! This I found utterly baffling.

In pondering how I might more effectively transmit my admittedly nonscientific, totally emotional concerns to a wide audience and gain their sympathy, I somehow came up with the idea of putting my ideas into rhyming quatrains. And so, before long, I had converted a great deal of the lecture into verse. As I tried it out on audiences, I found that the serious ideas in my message, now “leaner and meaner,” seemed to reach more people. Perhaps part of the reason for this is that I put on a kind of artistic persona in my rhymes, which allowed me to express myself in a more personal manner than I would dare to do in prose.

The first time I gave my versified lecture was, amusingly, in a back-to-back pair of lectures with Dave Cope right on his home turf in Santa Cruz, and our complementary talks went over very well. I might add that Dave himself—as one might expect, since music is his profession—belongs to that first category (the most intense lovers of music), and he and I even share a great deal in musical taste. This makes the discrepancy in our attitudes toward Emmy all the more striking, and, needless to say, thought-provoking.

Is Music Just Splicings of Licks, and No More?

Without further ado, let me now proceed to describe Emmy a little bit, and then begin giving my reactions in verse form. The basic idea behind Emmy is what Dave Cope terms “recombinant music”—the identification of recurrent structures of various sorts in a composer's output, and the reusing of those structures in new arrangements, so as to construct a new piece “in the same style.” One can thus imagine feeding in Beethoven's nine symphonies, and Emmy coming out with Beethoven's Tenth (or Brahms' First, if you subscribe to the claims of some musicologists that in his First Symphony, Brahms carried on the Beethoven spirit beyond the grave).

Toward the beginning of *Computers and Musical Style*, his first book about Emmy, Cope says this about his personal pathway of exploration:

In 1981, during a moment of recklessness, I wrote the following in a daily journal:

I envision a time in which new works will be convincingly composed in the styles of composers long dead. These will be commonplace and, while never as good as the originals, they will be exciting, entertaining, and interesting. Musicians and non-musicians alike will interact with programs that allow them to endlessly tinker with the styles of the composing programs . . . I see none of this as problematic. Machines, after all, only add and subtract. Programs that benefit from those operations are only as good as their creators.

This book describes many aspects of a program I have since devised for the replication of musical styles . . . If there is a discovery here, it is that one way of defining style is through pattern recognition and that musical style can be imitated if one can find what constitutes musical patterns. (Cope 1991a, p. xiii)

Here, then, is my opening salvo of quatrains in reaction to Cope's characterization of musical style as patterns.

Is music a craft,
Or is it an art?
Does it come from mere training,
Or spring from the heart?

Is music just notes,
Merely patterns combined
By a cocktail-bar pianist
With a wandering mind?

Though Fats Waller's ticklin'
Suggests profound joy,
Might it all be illusion
From a practiced riff-boy?

Does music, like poetry,
Cry from one's core,
Or is it just splicings
Of licks, and no more?

Do the études by Chopin
Reveal his soul's mood,
Or was Frédéric Chopin
Just some slick "pattern dude"?

Was Chopin a zombie with
The gift of piano gab?
Did he toss off mazurkas
Much as party bores blab?

Could he turn off his brain
And continue to sing
In true heart-rending fashion—
Or would one miss some zing?

Was Bach a musician
 Or mere *Musikant*?
 Did Johann his passion
 Express—or just cant?

In the furnace of Bach,
 Did there burn a pilot light,
 Or did Joh. Seb. compose
 On cool autopilot flight?

There's music that's trite,
 And there's music that's deep—
 Or is that the truth?
 Does all music come cheap?

Can one bypass the soul,
 Can one sidestep all strife,
 And produce wondrous music
 Without living life?

That's the crux of my talk;
 The idea, I hope, 's clear.
 And until recently,
 I myself had no fear.

A skeptic shot through,
 But then one day I heard
 Some not half-bad tunes
 From a program. My word!

So can style be learned
 By mechanical means?
 Can Rodgers be churned
 Out by Hart-less machines?

Soul-fire in Cole Porter
 Began his Beguine;
 Can we order more Porter
 From a Cole-less machine?

Well, so begins my commentary—making no bones about setting forth an emotional point of view. But if one is to form an educated opinion of Emmy, one's first

duty is obviously to familiarize oneself with how the program works. Cope, naturally, has his own ways of explaining Emmy, but I have found it useful to rephrase what I have learned over these past few years, and I think that hearing it from an outsider’s viewpoint may help to clarify certain difficult points. Moreover, I found more than once, in talking with Dave, that he would provide highly revelatory answers to key questions—questions that were not answered anywhere in his writings, and in fact in most cases were not even posed in his books. Such interchanges gave me a kind of personal insight into some aspects of Emmy that I believe may be useful to share, and so, with that as my excuse, I now present my amateur’s capsule portrait of Emmy’s innards.

A Personal View of How Emmy Works

Emmy’s central modus operandi, given a set of input pieces (usually all by a single composer and belonging to the same general form, such as *mazurka*) is:

(1) chop up; (2) reassemble.

This, in three words, is what Cope means by the phrase “recombinant music.” Caveat: The assembly phase, in contrast to Mozart’s famous *Musikalisches Würfelspiel*, which produced waltzes by random shuffling of 3/4 measures, is anything but haphazard or willy-nilly (as if by throwing dice). There are significant principles constraining what can be tacked onto what, and these principles are formulated so as to guarantee coherence (at least to the extent that the input pieces themselves are coherent!). I summarize these two principles as follows:

1. Make the *local flow-pattern* of each voice similar to that in source pieces.
2. Make the *global positioning* of fragments similar to that in source pieces.

These could be likened to two types of constraints that a jigsaw-puzzle solver naturally exploits when putting together a jigsaw puzzle:

1. The *shape* of each piece meshes tightly with those of neighboring pieces.
2. The *stuff* shown on each piece makes sense in the context of the picture.

The former of these constraints might be characterized as *syntactic meshing*, or meshing based solely on *form*, while the latter could be characterized as *semantic meshing*, or meshing based solely on *content*. In isolation, perhaps neither of them would be too impressive, but when used together, they form a powerful pair of con-

straints. But how does my jigsaw-puzzle metaphor translate into specific musical terms?

Syntactic Meshing in Emmy: Voice-Hooking and Texture-Matching

Let me first consider the first of these constraints—that involving form, or what one might call “coherence of flow.” This constraint in fact breaks down into two facets:

(1) voice-hooking; (2) texture-matching.

To understand these two distinct facets of syntactic meshing, one has to imagine that a new piece is being put together note by note, in sequence, and that to this end, short fragments of input pieces are being selected so as to mesh with the current context. Imagine that we have just inserted a fragment f_1 , and are considering whether to insert fragment f_2 right after it, drawn from somewhere in the input. Voice-hooking would be the requirement that *the initial note of the melodic line of fragment f_2 should coincide with the next melodic note to which fragment f_1 led in its original context.* In other words, a given fragment’s melodic line should link up smoothly with the melodic line of its successor fragment. This is very much like saying that two puzzle pieces should fit together physically.

Of course, here I referred only to the melodic, or soprano, line of a piece. One can also insist on voice-hooking of the bass-line, and of intermediate lines as well (tenor, alto, and so on). Ideally, voice-hooking can be carried out successfully on all voices at once, but if not, then the most logical voices to sacrifice are the inner ones, then the bass-line, and last of all, the melodic line. Usually, provided there is a sufficient quantity of input pieces, it will be possible to achieve a good deal of satisfaction in voice-hooking.

In addition, there is *texture-matching*, which is basically the idea that *the notes in a chord can be moved up or down pitchwise by full octaves and can be spread out time-wise so as to match some preexistent local pattern in the piece being composed.* Most typically, these two operations result in the “spinning-out” of a simple chord into an arpeggio that matches some preestablished arpeggiation pattern. Thus, a purely vertical C–E–G triad could be spun out, for instance, into a C–G–E–G figure to be incorporated into an Alberti-type bass-line, or into a very wide E–C–G arpeggio to match the widely arpeggiated pattern of the bass-line of a Chopin-like nocturne. It could even be turned into the very long sequence of notes “C–E–G–C–E–G–C–E; C–E–G–C–E–G–C–E,” which you may recognize as the melody in the first measure of the C major Prelude of Book I of Bach’s *Well-Tempered Clavier*. Basically, the

pattern of that piece is so regular that it is a mechanical act to spin out a triad into a whole sixteen-note sequence.

Semantic Meshing in Emmy: Tension–Resolution Logic and SPEAC Labels

We now turn to the second constraint—that involving content, or what one might call “tension–resolution logic.” This is where ideas devised by Cope as part of Emmy may in fact constitute a significant new contribution to music theory. The basic idea is that one wishes to insert a fragment into a new piece only if *the “location” of the insertion is similar to the “location” of the fragment where it occurred in some input piece*. The word “location” is put in quotes here because it is not clear what it means. Indeed, the italicized phrase forces one to ask the puzzling question, “How can a given fragment be ‘in the same location’ with respect to two different pieces? How can one compare ‘locations’ inside totally different pieces? What, indeed, might ‘location’ inside a piece be taken to mean (since, self-evidently, using measure number would be a pathetic travesty of an answer)?”

Cope decided that “location” must be defined in a way that involves both global and local contexts—in fact, a series of nested contexts, ranging from very local (notes, measures), to medium-range (phrases), to large-scale (periods), to global (sections). To a fragment on any of these distinct hierarchical levels (and there can be any number of such structural levels), Cope attaches a label—one of the five letters *S, P, E, A, C*—which attempts to capture what I have chosen to call the *tension–resolution status* of that fragment. These letters stand for the following words: *statement, preparation, extension, antecedent, consequent*. The label-assignment process proceeds from most local to most global, with the labels of larger sections dependent upon the labels already assigned to their component pieces.

Unfortunately, the details of the label-assignment process are unclear to me, but in essence it starts at the most local level, where the presence of specific scale degrees in the various voices is used as the main diagnostic for the labeling of a chord (co-presence of tonic and dominant, for instance, or tonic and mediant, suggests an *S* label at that level). From there on out, certain characteristic sequences of local labels are telltale cues that suggest specific higher-level labels, and so on, always moving upward hierarchically. In the end one winds up with SPEAC labels attached to sections of many different sizes and, perforce, at many different structural levels.

The upshot of this many-leveled labeling process carried out by Emmy is that any local fragment of an input piece winds up with a set of labels—its own label, that of the larger fragment inside which it sits, then that of the next-larger fragment in which

that one sits, and so on, and so on. Thus hypothetically, a given chord in an input piece could have the following set of labels (proceeding from most local to most global): A–C–C–E–P–A–S, and another chord might have the hierarchical series of labels E–S–C–S, and so on. In either case, such a series of letters basically tells you, on several different hierarchical levels, just what the tension–resolution status of the piece is at the chord concerned. And that—provided it really works well—would seem about as good a way of saying “where you are” in a piece as any I could imagine, since tension and resolution on many levels really do constitute the crux of musical meaning.

Now the trick is to use these labels to guide composition, and the basic idea is fairly straightforward. Suppose that in our piece-under-construction we find ourselves in a location whose tension–resolution status is PACSCS (moving from most local to most global). The letters *P-A-C-S-C-S* tell us “where we are,” so to speak, inside our new piece. And so, in choosing a fragment to borrow from an input piece and to insert right here, our main criterion will naturally be that the chosen fragment’s tension–resolution status inside its original piece was exactly PACSCS—in other words, that the fragment we are going to quote lies in “the same place” inside its original piece as in the new piece.

If in the input corpus we find several such “same-location” fragments, that is good, since it gives us a choice of how to continue, but we of course also want to satisfy the syntactic voice-hooking constraint. We thus throw away any fragments that do not match in this manner. If after this paring-down, there are still several potential fragments surviving and vying with each other for insertion, then we can choose one at random.

Suppose, on the other hand, that there is no input fragment that has exactly the desired multilevel tension–resolution status—how then to proceed? The only solution is to sacrifice something—but what? Cope decided that in such circumstances, global status is more sacrificeable than local, and so we lop off the final letter, leaving us with PACSC, and now we try again to find an appropriate fragment in the input corpus. If this fails, we lop off one more letter (thus giving PACS), and we search again in the input corpus. Since through such lopping-off we are loosening ever further the constraint of matching tension–resolution status, we will eventually find one or more input fragments that match the labels that we seek, and then we can choose randomly among those fragments, provided that voice-hooking also works. And thus the piece gets extended a little bit. At this point, we restart the constrained search process and extend the growing composition a little bit more—and so forth and so on. Thus, like a crystal growing outward, is built up a piece of music by Emmy.

In summary, here, in my own words, is the core of Emmy's composition process:

Sequential assembly of fragments that have the highest possible degree of agreement of SPEAC labels on all hierarchical levels

Stitching-together of fragments so as to respect voice-hooking constraints and so as to match local textures

Signatures

The preceding is the true core of Emmy, but in addition there are two other important mechanisms that should be described here as well. The first is what Cope calls *signatures*. A signature is a characteristic intervallic pattern that recurs throughout a composer's oeuvre, the use of which lends a high degree of seeming authenticity to a freshly composed piece. To find signatures, Cope has Emmy scour all input pieces for pairs of short note-sequences (say, between four and twelve notes, although there is no strict cutoff) whose intervallic patterns match either exactly or approximately. Thus, for instance, C–B–C–G would *exactly* match F–E–F–C, and would be a *near* match for D–C–D–A (the difference being that the first and second intervals are semitones in C–B–C–G, and whole-tones in D–C–D–A). Emmy scours the input for exact matches, and then gradually loosens up the search (relaxing the criteria governing interval-matching), until a satisfactory number of recurrent patterns have been found.

The variable numerical parameters in the computer code that determine whether a potential match is judged satisfactory or not are called “controllers,” and during a search for signatures, one must adjust the controllers until just the right number of signatures is found—not too few but not too many either. I know that in the past, Cope tended to do this adjustment of controllers himself in order to increase the effectiveness of Emmy's search for signatures, but perhaps by now he has managed to automate that aspect of the process. In any case, among the subtlest of controllers are those that winnow “insignificant” notes out of a given passage, leaving just “significant” ones; thanks to such controllers, Emmy can then match a highly embellished melodic fragment that contains, say, twenty very rapid notes with another melodic fragment that contains only four slow notes, and can discover the core signature that they share. Thus signatures found by Emmy can be very subtle indeed.

An important point is that such matching of intervallic patterns must take place *across* pieces, rather than *within* a given piece—for the obvious reason that any given piece will reuse its own motives many times, and Cope is not trying—indeed, he does not wish—to get Emmy to reproduce the melodic lines of a given piece, but rather he

wishes Emmy to pick up on and to exploit the recurrent (but less obvious) melodic patterns that a composer tends to reuse from piece to piece, probably without even being aware of doing so.

It may not seem a priori evident, needless to say, that all composers do have signature motives, but this has turned out to be the case. One might tend to think that the existence of many signatures would show that a composer is rut-bound, and perhaps it does, but in any case, it is a universal fact, revealed in undeniable fashion by Cope's work on Emmy, that each composer does employ interval-pattern motives that recur in piece after piece.

Once such signatures have been identified in the input, they are stored in a database, with each diverse instance of a given signature being stored *together with its underlying harmonies*, thus all ready for insertion *as a whole* inside a new piece. You might suppose that the insertion of prepackaged, precisely quoted chunks would risk producing passages that sound like pure plagiarism, but surprisingly, these prepackaged chunks are usually so generic-seeming and so small that, even to a highly astute listener, they don't shout from the rooftops which precise piece they came from; they merely sound like the given composer in a nonspecific, nonpinpointable manner.

Templagiarism

The second mechanism that I wish to describe here is what I dub "templagiarism," short for "template plagiarism"—a fascinating, more abstract version of the signature concept. If, in scanning a given input piece, Emmy notes that a motive appears in quick succession two or more times (again with some liberty taken in the matching, thus allowing variants of a given motive, such as tonal entries of a fugue theme, to be counted as "equal" to each other), it records the following data for these entries: (1) the *pitch displacement* of the new occurrence relative to the previous occurrence, and (2) the *temporal displacement* of the new occurrence relative to the previous occurrence. In short, Emmy records, for any repeated motive, the "where-and-when" pattern that characterizes the motive's repetitions. Emmy then detaches this abstract pattern from the specific motive in question, and takes it to be characteristic of the composer's style. Note that this is a higher-order architectural stylistic feature than a mere signature, because it is concerned not with any motive itself but with how that motive recurs within a piece.

Templagiarism can be an astonishingly effective style-evoking device, as I found out one day when listening, in Cope's living room, to "Prokofiev's Tenth Sonata for

Piano” (as Dave humorously, or perhaps hubristically, dubs one of Emmy’s pieces, about which more later). As the second movement started, I heard a very striking chromatically descending eight-note motive in midrange, then moments later heard the same motive way up high on the keyboard, then once again a few notes lower, and then one last time very deep down in the bass-line. These widely spread entries gave an amazing feeling of coherence to the music. Indeed, for me the passage reeked of Prokofievian impishness, and I thought, “Good God, how in the world did Emmy do *that*?” It sounded so well calculated (not in the computer sense of the term!), so inventive, so full of musical intelligence.

Astonished, I asked Dave what was going on, and he replied, “Well, somewhere in one of the input movements on which this movement is drawing, there must be some motive—totally different from *this* motive, of course!—that occurs four times in rapid succession with exactly these same timing displacements and pitch displacements.” Then he spelled out more explicitly the concept of templagiarism to me. It would have been pleasing if at that point we had scoured Prokofiev’s scores until we found exactly such an episode, but we didn’t take the trouble to do so. I’ll take Dave’s word for it that we would find it somewhere or other.

Cope’s idea of templagiarism is itself brilliant and devilishly impish: it borrows a touch of genius from the composer at such a high level of abstraction that when the pattern is simply quoted lock, stock, and barrel—plagiarized, no more, no less—it once again sounds like a touch of genius, but an utterly fresh and new one. The reason it sounds fresh and new is, of course, that in order to quote the template, you need to supplement it with a new “low-level” ingredient—a new motive—and so the quotation, though exact on the *template* level, sounds truly novel on the *note* level, even if one is intimately familiar with the input piece from which the template was drawn. New filler material has been spliced into an old template that bears the unmistakable stamp of a specific genius, and so the whole passage has a powerfully compelling feel to it—a deep musical mind seems to lie behind it.

It’s a bit as if one were to use fancy speech-synthesis technology to make the very familiar voice and accent of, say, John Kennedy come out with utterances that Kennedy himself never made—perhaps nonsense statements, perhaps cheap rabble-rousing inanities that he would have despised, whatever. Despite their falsified content, they would still sound for all the world like Kennedy (at the voice level, at least), and such statements probably would seem genuine to most people.

I must admit that I don’t have a clear understanding of how the very complex operation of templagiarism (or, for that matter, the somewhat simpler operation of insertion of signatures) is made to coexist harmoniously with the previously described syntactic and semantic meshing-operations, because I can easily imagine them con-

flicting with each other. Nor do I understand how Emmy composes a “motive” and deems it worthy of use as such in an extended movement. But of course, how could I? It would probably take many months of intense study of Emmy to understand such matters. I remain an intrigued outsider, and hope and expect that over time, Dave will explain Emmy’s principles ever more lucidly.

The Acid Test: Hearing and Voting

The foregoing provides a summary of what I myself have absorbed about the workings of Emmy, both from reading Cope’s books and from a good number of one-on-one conversations with him. We now continue with a few more of my quatrains about Emmy.

David Cope, a composer
At UCSC,
Has a program make music
From S, P, E, A, C.

Cope’s “EMI” takes scores
By, say, Bach—scores of scores!
Then it scours these scores
For Bach-style “signatures.”

From a “style-free” scaffolding
(A pattern of “SPEAC” ’s),
The program hangs signatures,
And lo! Old Bach speaks!

So is music an art,
Or is it merely a craft?
Remember at whom it was
That they all laughed.

The proof’s in the pudding
(In this case, the ears);
If you’ve not heard EMI,
Don’t prejudge it with sneers.

At this juncture in my lecture, I have almost always had a live pianist—sometimes Dave’s wife Mary Jane Cope, who is on the music faculty at UC Santa Cruz—perform

a handful of small two-voice pieces for the audience. The listeners are forewarned that there is at least one piece by Johann Sebastian Bach in the group, and at least one by Emmy in the style of Johann Sebastian Bach, and they should try to figure out which ones are by whom (or by what).

As a prelude and to set the proper tone, I first read aloud the following two short excerpts from Cope's *Computers and Music Style* (Cope 1991a), the first one describing a very simplified version of Emmy which Cope devised solely for pedagogical purposes, and the second one ushering in the chapter in which the full-strength Emmy—at least the Emmy of that vintage—is carefully discussed (though it is certainly not described in full detail):

It will create small two-part inventions similar in nature (not in quality) to those created by Bach. (p. 98)

For the true inheritance of Bach's style to take place, a much more elaborate program would be necessary. This more elaborate program is presented in the description of Emmy in the next chapter. (p. 136)

Make of that telling little phrase “the true inheritance” what you will . . .

After the pieces have been performed, I tell the audience that they are now going to vote (with the proviso that anyone who has recognized a piece from their knowledge of the classical repertoire is disenfranchised). The result has usually been that most of the audience picks the genuine Bach as genuine, but usually it is only about a two-thirds majority, with roughly one third getting it wrong. And it is not by any means always the less sophisticated audience members who make the wrong classification. In any case, once people have made their vote, I then return to my verse, as follows:

Well, now you've heard EMI,
Perhaps you feel had.
In your shoes, so would I.
When one's fooled, one feels bad.

And if you were right,
Not a single guess wrong,
You've the right to feel smug
For a while—but how long?

So you told Bach from EMI,
So you've got quite keen ears;
But EMI's evolving—
Just wait a few years.

To Sound like Bach and to Speak like Bach

It is indeed true that Emmy is evolving—it is a moving target. Cope began work on his program in 1981, and in all these years he has not let up on it. Emmy's early pieces are, like any fledgling composer's, pretty amateurish affairs, but her later output sounds increasingly impressive, and Cope has grown more and more ambitious over time. Whereas initially he was proud of Emmy's production of short two-part inventions and short mazurkas, he now has Emmy producing entire sonatas, concertos, and symphonies. There is even a "Mahler opera" under way or in the works—something that would certainly be a challenge for any human composer to carry off.

What exactly is the difference between stylistic imitation as carried out by a human being and stylistic imitation carried out by a computer program? My friend Bernard Greenberg has been writing music in the style of J. S. Bach (and other composers, but Bach most of all) for decades. Indeed, among my life's most amazing memories are of visits to Bernie's apartment, where, as I listened to him play his own soulful pieces on the organ, filled with complex dissonances and marvelously unexpected turns of phrase, I felt as if I were in the presence of Old Bach himself. One time I brought along a mutual friend to listen, and he—also a gifted musician—made the following unforgettable remark to Bernie: "Gee, not only is your music in the Bach style but it *sounds* good, too!" I always found this remark extremely puzzling, since to me the very *essence* of Bach style is that it "sounds good." How could something possibly sound deeply *Bach-like* and yet also sound *bad*? The tone of the remark made no sense to me—and yet I must admit that Bernie himself once made a related remark about the secrets of capturing Bach's style: "The trick is to make music not that *sounds* like him, but that also *speaks* like him."

The Nested Circles of Style

Well, of course, what is being hinted at here, though in a blurry way, is that style is a multilayered phenomenon. There are shallow aspects to style (how a piece "sounds," in Bernie's terms), and then there are deep aspects (how it "speaks"). It is quite possible that someone could be capable of capturing many of the shallower trademarks of a composer and yet miss the bull's-eye as far as essence is concerned. I always think of Schumann's short piano piece called "Chopin," which occurs in his *Carnaval*, which on one level "sounds like" a Chopin nocturne—it has the characteristic wide left-hand arpeggios and a lot of melodic embellishment—and yet on a deeper level it quite misses the mark in terms of Chopin soul (at least to my ear).




Figure 2.1

An extremely simple yet seemingly inevitable diagram pertaining to stylistic imitation.

This talk of different levels of style and of targets and bull’s-eyes suggests the following extremely simple yet seemingly inevitable diagram pertaining to stylistic imitation (see figure 2.1).

Someone who glibly captures only the most obvious features of a composer’s style—an Alberti bass, say, for Mozart—would fall in the outer ring but leave all inner rings untouched. A deeper imitator would add other outer layers of style but fail to penetrate all the way to the core, or stylistic bull’s-eye. But only someone who had dedicated years to the art, and whose emotional makeup, moreover, bore a deep affinity to that of the composer in question (and this is how I see Bernie vis-à-vis Bach), could hope to come close to that elusive central core that constitutes true Chopinity or Bachitude.

And yet . . . there is something most troubling to me about this diagram, as I have drawn it—namely, the fact that the ring with the greatest area is the outermost one, not the innermost one. This disturbs me because it suggests that you will get the most effect from the simplest and shallowest tricks. The diagram suggests that as you proceed further and further in—as your mastery of the art ever deepens—the area you are adding becomes smaller and smaller. When you have acquired but one layer of style mastery, your music will surely not fool experts, but it might fool 80 percent of the general populace. Work harder, add the second ring of mastery, and now you fool 90 percent. Add the third ring, and your fooling rate goes up to, say, 95 percent, and the fourth ring gets you to 98 percent. There’s still something missing, but sadly, the missing ingredient is getting subtler and subtler, tinier and tinier . . . In the end, then, with all but the innermost circle, you may wind up reliably fooling all of the

world's top experts, while still lacking Bach's true soul. In short, it's a most depressing thought, if the nested-circles image is accurate, that the innermost layer, though surely the most difficult of all layers to acquire, is also the smallest and perhaps, therefore, the least significant in terms of its effect upon listeners.

There are layers of style
 From the skin to the core.
 The former are patterns;
 The latter—something more?

If style's many layers
 Are like circles that nest,
 Then the ones near the crux
 Grow more tiny. I'm depressed.

When Does a Beatles Song Sound like a Bach Chorale?

In an e-mail exchange with me, Bernie Greenberg was discussing his attempts to impart to others his ability to write Bach-like music, and he wrote this:

There are tricks of the trade, and you can teach chorale-writing such that anyone with a little talent can write a chorale that sounds like a Bach chorale *that you are not listening to closely*.

A little later in that same e-mail exchange, in relating an episode in which he had helped an acquaintance who wrote four-part chorales and who wanted Bernie's advice on how to get them to sound more Bach-like, Bernie amplified his remarks as follows:

There is no question that by further refinement of style, I can make them sound more like Bach chorales than many other church hymns. Perhaps the right question is:

“Do they sound more like Bach chorales than *what?*”

rather than

“Do they sound like Bach chorales?”

After all, compared to jet takeoff noise, or even Balinese gamelan music, most Beatles songs “sound like Bach chorales,” right?

A Portrait That “Looks like” Its Intended Subject

Bernie's humorous point is right on the mark, and forces one to think carefully about what it means to say glibly, *X sounds like Y*. And further light is shed on the question by considering the analogous issue of what it means to say, *X looks like Y*. To make



Figure 2.2
A standard “smiley face” image.




Figure 2.3
A few strategically placed parallel vertical lines added to the bland, generic smiley face, demonstrating a tiny amount of “style.”

this issue vivid, let us take a standard “smiley face” image, as shown in figure 2.2. Presumably, the bland face shown below does not remind you of any individual you know, right?

It would be surprising if it did. But if we now add to our bland, generic smiley face a tiny amount of “style”—just a few strategically placed parallel vertical lines—lo and behold figure 2.3!

All of a sudden, nearly everybody recognizes the familiar face of the *Führer* of the Third Reich. To be sure, nobody would say about this mustachioed inanity, “It looks very much like Hitler”; perhaps nobody would even say, “It looks like Hitler”; but despite that, everybody *sees* Hitler in it. They can’t help it. The point of this example, invented by David Moser (who grimly subtitled his ironic image “Have a Nice Holocaust!”), is that just a minimal gesture in the direction of a known style can, if well executed, have a stunning effect, summoning up much more than is really there.

So . . . how much are we being fooled when, on hearing a piece of music, we respond to some gestures that in the past we have come to associate with composer X, and then exclaim to ourselves, “This piece sounds like X?” Can we even distinguish clearly between responses at a shallow level and a deep level? Indeed, what *is* the difference, in music, between “shallow” levels and “deep” levels of style? Is it just a question of different levels of depth of syntactic pattern, or is it something more than that?

Lewis Rowell's "Bach Grammar"

Not long after I became a professor at Indiana University, I heard on the radio a very engaging piece for organ that to my ear sounded extremely Bach-like; when it was announced, however, I found out to my surprise, though not to my chagrin, that it had been composed by a music professor at IU—Lewis Rowell. I lost no time in contacting Rowell and suggested we have lunch together to talk over the idea of faking Bach. He was delighted that someone had taken an interest in his piece, and we soon met. Over lunch, I asked Rowell how he had composed such an authentic-sounding piece, and he said, "Oh, that's not hard . . . Bach developed a kind of grammar that I merely picked up, as could anyone who wished to. And then, armed with this grammar, I—just like anyone with a bit of musical talent—can easily compose any number of pieces in perfect Bach style. It takes no genius, believe me. It's all straightforward stuff. The only place where genius was involved was in coming up with the grammar."

I was astounded to hear how dismissively Rowell described his acquisition of "Bach grammar," and just as astounded to hear that he thought that composing long, complex, and coherent new pieces in the full Bach style was basically merely a mechanical act, requiring no act of genius whatsoever. After all, I, a lifelong lover of Bach, had on several occasions tried composing pieces in the Bach style, and had found myself unbelievably stymied. Measures and short phrases, yes, perhaps—but a long movement? No way!

Rowell's claim, however, was that only Bach's own *creating* of his supposed "grammar" was hard, whereas *inducing* that grammar from Bach's output and then *exploiting* it was a piece of cake. A glib hack could create new works as deep and as great as any that had ever issued from the pen of the great Baroque master—or from that of any other great master. Profundity becomes a snap, emerging at the drop of a hat. By contrast, my personal feeling, based on my own experience (and, I must say, based also on long observation of Bernie Greenberg), was that *extracting* a true and deep "Bach grammar" from Bach notes was itself an act that would require extraordinary insight—perhaps even genius. And even if such a grammar could be extracted (which struck me as highly implausible, Rowell's claims notwithstanding), I felt that to *exploit* it to make new pieces as great and as engaging as those of J.S.B. himself would still be an act of enormous creativity.

Many years later, grappling mightily with the strange new world of Emmy and her algorithmically induced grammars, I remembered my stimulating lunch with Lew Rowell and wondered what Bernie Greenberg would think of our chat. So I sent

Bernie the gist of Rowell's claims through e-mail, to which he quickly responded with the following eloquent set of remarks in his own inimitable style (if I dare make such a claim!):

I'd be very interested in such a grammar. It would have to include a "syntactic description" of the theology of Paul as expressed in Romans, the innate intellectual tension between Christ's roles as Victim and Savior, and other emotional basis vectors of the space which is "Bach."

Anyone who has been moved by the *St. John Passion*, the *St. Matthew Passion*, or the Cross dialogue of Cantata 159 understands that the root of their emotional power is in the turgid psychodynamics of the Crucifixion, not in the seventh-chords, which are the mere paint that Bach has used to implement these canvases, incomparable paint though it be.

Although I sympathized with what Bernie was trying to say, I felt he had overstated the case. Does one really need to be a pious Christian to be able to compose deeply Bach-like music, or even to be powerfully moved by Bach's music? In point of fact, Bernie himself, brought up Jewish and an atheist by credo, provided a counter-example. I argued that the essence of Bach's power comes not from his deep piety but from his deep humanity—from just those human experiences discussed in my speculation (quoted from *GEB*, above) about a computational "music box" producing new Bach and Chopin pieces. Bernie, on hearing this objection, conceded that among the most important "emotional basis vectors of the space which is 'Bach'" are many that have nothing per se to do with religion but that simply follow from being born into this crazy world, growing up in it, and living a full human life. And so Bernie closed his musings by saying this:

When the "grammar" is sufficient to cover such notions, the General AI problem will have been solved, I think.

Amen. As for myself, I was inspired by all these musings on alleged "Bach grammars" and the hidden inner fire of human creativity to write the following series of quatrains.

When music's been treated
By the likes of Dave Cope,
Is the mystery banished,
Or is there still hope?

Does true depth in music
Mean creating new styles,
So that music by mimics
Is worth just snide smiles?

Was Chopin's fourth ballade
 A mere splicing of licks
 From his previous three—
 Or were there new tricks?

What's creative? What's rut-stuck?
 What is new, and what's old?
 What's derivative? What's novel?
 What is weak, and what's bold?

Is a style, once devised,
 A mere snap to ad lib
 A bunch of new tunes in,
 Provided you're glib?

Is Bach-style a grammar
 A hack can acquire,
 Or is there some essence—
 Some deep inner fire?

Just what makes a genius
 Than a mimic far better?
 The former forges spirit;
 The latter worships letter.

'Twixt genius and mimic,
 What makes the sharp cut?
 The former's unfettered,
 The latter's in a rut.

Showing Up Despite Being a No-Show

When I gave my talk with Dave Cope at Santa Cruz in May of 1996, I was hoping to persuade Cope's Santa Cruz colleague Tom Lehrer, of satiric-song fame, to take part in a panel discussion on Emmy, and to that end, I called up Lehrer (whom I had known for some years and who I knew was something of a recluse) and tried to persuade him to join us. He was, however, predictably self-deprecating and in the end turned me down, although in the nicest of ways. In fact, our phone chat lasted at least ten to fifteen minutes, and I found what he said very provocative. When I hung up, I all of a sudden realized that although Lehrer had declined to come, he had actually

told me over the phone pretty much what I had most hoped he might say in front of a live audience. Given that irony, I quickly jotted down everything that I could remember, which was a lot, and then promptly translated it into verse. I figured that this way I could give my audience a “virtual Lehrer” (and keep in mind that the German word *Lehrer* means “teacher”) if not the real McCoy.

I must admit that I was also secretly hoping that Lehrer would show up in person at the back-to-back lectures Dave and I were giving, because then I could play my little joke on him, of surprising him by delivering his own ideas, in versified form, to the assembled group despite his having declined to participate. Unfortunately, no such luck—Lehrer didn’t attend our talks. Nonetheless, my Lehrer quatrains were appreciated by the audience, and I feel they enrich the whole discussion, and so, for what they’re worth, here they are:

A teacher I know
Whom I asked to take part
In this meeting, said, “No,
What I do is no art . . .

“I’ve nothing to tell folks;
I won’t take the stand.
It’s true, I write songs,
But they’re boring and bland . . .

“You just name me a form
Such as ‘march,’ and I’ll play
You a piece with a march beat,
Cliché after cliché . . .

“All my songs are deriv—
They’re in nobody’s style.
If I try copying Kern,
It comes out sounding Weill!

“Still, old Irving Berlin
Has a style I might snag,
For his music’s as patterned
As a Scott Joplin rag . . .

“Berlin plays vanilla
To Kern’s chocolate mint;
So I might stamp out tunes
From that old Berlin mint . . .

“But it’s truly a cause for despair
When you come to the genius of Kern.
He pulls magical chords from the air
With an ease too profound to discern . . .

“Oh, I guess if I truly did yearn
To mimic the magic of Kern,
I could study and probe and might learn
Some tricks that make Kern phrases turn . . .

“Then armed with this kernel of Jerome,
I might slowly begin out to churn
The patterns that once seemed so special—
The signatures of the great Kern . . .

“But even at that advanced stage,
Pulling wool over Kern experts’ eyes,
The flame of the novel I’d lack—
Lacking genius, I’d just plagiarize . . .

“No, the greatness of Kern I can’t ape;
He’s a doctor, I’m merely a quack.
And that’s why I wouldn’t belong
On a stage with you folks with the knack!”

Ah, the irony of his remarks!
Here’s a fellow who *does* have the knack
To spin songs in the styles of yore,
Yet declines, saying, “I’m just a hack!”

Yet my friend in declining said so much,
And so well that I wrote it all down,
Then converted it into this verse,
So he’s here despite turning me down!

His modesty struck me as odd.
Just why, if the Kern style did yield
Its keys to his scan, would he say,
“Still I’m nought in the novelty field?”

To what higher goal could one aspire,
Than the crafting of tunes on one’s forge—
Be they Kern-style, or Lerner & Loewe,
Or Bernstein, or Lehrer—or George?

Lennie Is Jealous of George

My allusion to “George” here is, specifically, a reference to George Gershwin. The reason for this is that I was deeply struck when I read, in Leonard Bernstein’s *The Joy of Music* (Bernstein 1959), an article provocatively called “Why You Don’t Just Run Upstairs and Write a Nice Gershwin Tune?” The article is in the form of a dialogue between L.B. himself and a character called P.M. (Professional Manager). The two of them are meeting over lunch, and we tune in on their conversation as they chat ‘n’ chew:

P.M.: Learn a little from George. Your songs are simply too arty, that’s all. George didn’t worry about all that. He wrote tunes, dozens of them, simple tunes that the world could sing and remember and want to sing again. You just have to learn to be simple, my boy.

L.B.: You think it’s simple to be simple? Not at all. I’ve tried hard for years. A few weeks ago a serious composer friend and I were talking about all this, and we got boiling mad about it. Why shouldn’t we be able to come up with a hit, we said? So we went to work with a will, vowing to make thousands by simply being simple-minded. We worked for an hour and then gave up in hysterical despair. Impossible. I remember that at one point we were trying like two children, one note at a time, to make a tune that didn’t even require any harmony, it would be that obvious. Impossible. It was a revealing experiment, I must say, even though it left us with a slightly doomed feeling.

Let me quote from the fellow whose lyre
Gave us *West Side Story* and *Candide*—
He’s a dragon whose music breathes fire,
Yet he sighed, “By George, I’m out-keyed!”

A fake luncheon chat he once penned:
“Hey, Why You Don’t Just Run Upstairs,
And Write Me a Nice Gershwin Tune?”
Its point was the depth of simple airs.

Len denies, in this chat,
That new tunes he could spin
That would capture the essence
Of his idol, Gershwin.

Indeed, his whole point
Is the fact that it’s tough—
Not just tough but damned tough—
To make new Gershwin-stuff.

You struggle and strive
 To be Georgishly alive,
 To be simple, to jive,
 Yet you never arrive.

There's a spirit inside
 That just won't show its face,
 Though you hear it inside
 Every note, graced with grace.

Lenny's right, I would say:
 To dream up "I Got Rhythm"
 Takes something beyond
 A pure pattern algorithm.

Devilishly Infectious Rubbish Spouted by the Orwellian Versificator

While we're on the topic of famous Georges, there is another George whose ideas are highly germane to our topic. I speak of George Orwell and his frightening novel *1984* (Orwell 1949). When I read it in high school, many nightmarish images haunted me, but there was one odd passage that came flashing back to me from far across the decades when one spring morning in 1996 I caught myself humming, to my own horror, a certain mazurka in the shower . . .

And the Ministry had not only to supply the multifarious needs of the Party, but also to repeat the whole operation at a lower level for the benefit of the proletariat. There was a whole chain of separate departments dealing with proletarian literature, music, drama, and entertainment generally. Here were produced rubbishy newspapers, containing nothing except sport, crime, and astrology, sensational five-cent novelettes, films oozing with sex, and sentimental songs which were composed entirely by mechanical means on a special kind of kaleidoscope known as a versificator . . .

Under the window somebody was singing. Winston peeped out, secure in the protection of the muslin curtain. The June sun was still high in the sky, and in the sun-filled court below, a monstrous woman, solid as a Norman pillar, with brawny red forearms and a sacking apron strapped about her middle, was stumping to and fro between a washtub and a clothesline, pegging out a series of square white things, which Winston recognized as babies' diapers. Whenever her mouth was not corked with clothes pegs, she was singing in a powerful contralto:

*It was only an 'opeless fancy,
 It passed like an Ipril dye,
 But a look an' a word an' the dreams they stirred,
 They 'ave stolen my 'eart awye!*

The tune had been haunting London for weeks past. It was one of countless similar songs published for the proles by a sub-section of the Music Department . . . But the woman sang so tunefully as to turn the dreadful rubbish into an almost pleasant sound . . .

She knew the whole driveling song by heart, it seemed. Her voice floated upward with the sweet summer air, very tuneful, charged with a sort of happy melancholy. One had the feeling that she would have been perfectly content if the June evening had been endless and the supply of clothes inexhaustible, to remain there for a thousand years, pegging out diapers and singing rubbish.

I didn't recall this passage word for word, but the overall image had stuck accurately in my mind for over three decades, and reading it again made me cringe just as I had back then, imagining the ragingly infectious power of the formulaic, mechanical junk-music issuing forth from the quaintly described "kaleidoscope known as a versificator." But now, having been sucked in myself by the kaleidosCope known as Emmy, I really had no choice but to write the following verses:

One mazurka by EMI
Has lodged, I confess,
In the grooves of my brain,
Causing shame and distress.

Like the proles in George Orwell's
Nineteen Eighty-Four,
I find myself humming
An emotionless score.

I feel shock and bemusement
And confusion, to boot:
Is this *rubbish* I've swallowed?
Am I that unastute?

I never did dream
I'd be mortified by
Merely humming some tune;
Now I eat humble pie.

After decades of sureness
That the pieces I hear
Are deep mirrors of passion,
Must I now reverse gear?

The Pea-Sized Creative Module Keeps on Truckin' . . .

Peg Brand, a friend and colleague in Indiana's philosophy department, sent me through campus mail a remarkable article by Amei Wallach that she'd read in the *New York Times* (Wallach 1995). Its subject was the American painter Willem de Kooning, whose mental health had, in the early 1980s, suffered a sharp decline. Indeed, by the middle of that decade he was in the fullest throes of Alzheimer's disease, and yet the article was all about a series of paintings that he had executed during that period of his life, and that, rather astonishingly, had garnered high praise from not a few art critics. The curator of a large retrospective exhibit of de Kooning's last set of paintings called them "among the most beautiful, sensual, and exuberant abstract works by any modern painter." And yet one has to remember that during this period, de Kooning often would paint the same painting over and over again, unable to remember having done it before, until one day someone convinced him to keep his most recent canvases right in front of him as he worked, so he would see them and thus be able to avoid incessantly repeating himself.

Psychiatrists and neurologists interviewed in the article stressed that advanced Alzheimer's victims can no longer carry out any kind of mental activity that involves maintaining coherence—for instance, though they can play golf physically, they have to ask, on each stroke, where their ball is. Writing a sensible novel—in fact, most of the time, even uttering a sensible nonroutine sentence—is out of the question, as is sustaining a chain of reasoning beyond a few seconds. It is thus definite cause for pause to find out that paintings hailed as great art were produced by a mind whose light was day by day growing fainter, a mind whose owner no longer recognized any other human's face, no matter how long known, a mind that seemed to do nothing but wander in vague, aimless circles—when not filling blank canvases with lines and colors.

There're lots of old-timers
 Who still can create;
 But those with Alzheimer's—
 Can *their* art be great?

De Kooning is brain-dead,
 He paints as in sleep;
 Yet critics acclaim him:
 "Great stuff—makes you weep!"

Suppose that Old Chopin
 Had lived to 89,
 Losing all of his memory
 As well as his mind.

Yet when he sat down
 To make up a fresh tune,
 His magical chords
 Soon made listeners swoon.

What survived in his brain
 Was the size of a pea—
 A module for composing
 Autonomously.

When Chopin wrote waltzes,
 Did he draw on all life,
 Or could some “waltz module”
 Be excised with a knife?

Is composing a narrow,
 Mechanical skill,
 So old geezers can compose
 Using minds that are nil?

Is music, like chess,
 A wee, hard-edged domain,
 Algorithmically handled
 By a pea-sized subbrain?

Emmy Tries Her Hand at Doing Chopin

At this point in my lecture, I usually have the second musical interlude, this time involving two or three mazurkas, at least one by Chopin, at least one by Emmy. Rather than describing what happens myself, I would like to quote here what one person who was in the audience of my most recent Emmy lecture at the University of Rochester wrote to me and Dave Cope afterward.

From: kala pierson <kpi@ibm.net> Mon Feb 1 19:00:12 1999
 To: howell@cats.ucsc.edu
 Subject: EMI's big day at U. of Rochester ...
 Cc: dughof@indiana.edu

Hi, David! I heard Douglas Hofstadter's EMI-demo at the U. of Rochester yesterday; and though you'll probably hear an account of it from him, I wanted to give you a report from the trenches too, since EMI made such a dramatic impression on us.

As you know, Eastman School of Music is part of U.R.; much of the audience was made up of theorists and composers from Eastman (I'm a composer). DH gave us three listening tests: Bach inventions, live; Bach arias, on video; & Chopin mazurkas, live. It was easy for most of the audience to tell EMI from Bach; there were a lot of knowing smirks among those around me during the not-Bach inventions. Okay, we concluded, those imitations are pretty remarkable on several levels but they just ain't the real thing, and we--Those In the Know--can tell.

When the pianist played the two "Chopin" mazurkas, we were similarly confident. The first mazurka had grace and charm, but not "true-Chopin" degrees of invention and large-scale fluidity; there were very local-level, "shallow"-feeling modulations--just the type, I reasoned, that a computer program would generate based on more sophisticated examples of the real thing. The second was clearly the genuine Chopin, with a lyrical melody; large-scale, graceful chromatic modulations; and a natural, balanced form.

Although DH told us that the vote on this piece looked like "about 50/50" from his point of view, there was a definite preference among the theory/comp corner of the audience. I voted real-Chopin for the second piece, as did most of my friends. When DH announced that the first was Chopin and the second was EMI, there was a collective gasp and an aftermath of what I can only describe as delighted horror. I've never seen so many theorists and composers shocked out of their smug complacency in one fell swoop (myself included)! It was truly a thing of beauty.

Cheers for now,
kala

"Truly a thing of beauty!" This is an amazingly refreshing and candid statement from someone at one of the most elite music schools in the United States. Perhaps only a student could have written it. But no, I take that back. There are professors who are just as honest, though certainly it is hard to swallow one's pride and admit having been taken in. For many, it would be tempting not to admit having been gulled, and to go around haughtily pronouncing Emmy's compositions to be cheap forgeries of no artistic merit whatsoever. Indeed, I have run into some professional

musicians who have done just that. But I personally have heard too much of Emmy's music and been fooled too often to be haughty, even though some of it is certainly very weak (thank God!).

Suppose we discovered
A pristine Volume III
Of the "Tempered Clavier,"
With the depth of JSB.

It makes a huge splash,
And musicians galore
Compete to perform it
In grand halls the world o'er.

It meets with reviews
That are tops. All agree
These are fugues without peer:
"Ach, it's Bach—only he!"

But for some, strange to say,
If Dave Cope were to spill
Bitter beans—"It's by EMI"—
Then its worth would be nil.

They'd retract all their praise,
No more sing its great powers,
For now it's just fool's gold—
Bouquets of fake flowers.

Musicians a-plenty
There are, who, if told
In advance, "It's by EMI,"
Will find flaws. Ain't that bold?

But I fear that it's not,
For it's after the fact.
Forewarned "It's the en-EMI!",
They so "bravely" attacked.

I find it more honest
If one's judgment remains
Unswerved when one learns
It's by chips, not by brains.

When the votes were taken at my Rochester lecture, someone called out, “How’d we do compared to Indiana?” (my own university also having one of the nation’s top music schools). Everyone laughed, especially when it was suggested that perhaps a new system for ranking music schools could be based on how few errors were made by faculty and students in telling pieces by Emmy from pieces by human composers.

One stunning lesson from my Rochester lecture (and indeed, from all of the times I’ve lectured on Emmy) is that people with deep musical gifts and decades of training can, on occasion, mistake an Emmy product for the genuine article. And remember—we are just embarking, we humans, on the pathway toward the realization of the dream of “preprogrammed mass-produced mail-order twenty-dollar desk-model music boxes”—those boxes on whose “sterile circuitry” I heaped so much scorn, back when I wrote *GEB*.

Where will we have gotten in twenty more years of hard work? In fifty? What will be the state of the art in 2084? Who, if anyone, will *still* be able to tell “the right stuff” from versificator rubbish? Who will know, who will care, who will loudly protest that the last (though tiniest) circle at the center of the style target has *still* not been reached (and may never be reached)? What will such nitpicky details matter, when new Bach and Chopin masterpieces applauded by all come gushing out of silicon circuitry at a rate faster than H₂O pours over the edge of Niagara? Will that wondrous new golden age of music not be “truly a thing of beauty”? Won’t it be sweet to swoon in a sea of synthetic sublimity?

If output from EMI
Fooled all but an elite,
To protest, “Crux is missing!”
Would ring quite effete.

When the “heart” that is missing
Is unmissed by most,
Then the essence that’s missing
Is a wisp of a ghost.

And this is my fear—
That what’s missing will shrink
To near zero, with time.
And then—what to think?

When music’s reduced
To the schemas of Cope,
Has the romance all vanished?
I would like to sing “Nope.”

Scarfig Down Spastroni Down in South Delabam'

Have you ever downed a plate of delicious *spastroni* swimming in *caravinese* sauce? Or consumed *pollitucciollo* with a side order of *pomostacchi*? Ever eaten a salad with *pomodorini* and *marboli*? If you answer yes, I'd say you're a brazen bluffer, for these words were produced by a very simple computer program which had been "fed" (for want of a better word!) the names of many Italian foods, such as *spaghetti*, *ravioli*, *lasagne*, *vongole*, *rigatoni*, *pomodori*, *fettuccine*, *linguine*, *vitello*, *pollo*, *mostaccioli*, and so forth. There is no such thing as *caravinese* sauce or a side order of *pomostacchi*, my friend! (On the other hand, by sheer luck, there really are *pomodorini*—cherry tomatoes—for the computer-generated word just happens to be the standard diminutive of *pomodori*, "tomatoes.")

This kind of program, first shown to me in the early 1960s by my friend Charles Brenner, and which originated at Bell Labs in the late 1940s, is based on imitating the frequencies of short letter groups—in this case, *trigrams*. Given a passage of input text, the method is based on the probability of a pair of letters (*sp*, for instance) being followed by various other characters (*a*, for instance, would be a likely follower, as opposed to *q*, which would be nonexistent).

To generate a piece of output text that mimics the trigram frequencies of the input text, you begin with any digram (two-letter piece) of the input text (*sp*, let's say). Then you look at *which* characters followed that digram in the input text, and *how many times* each of them did so. Perhaps the letter *a* followed *sp* four times, *i* followed it once, *o* followed it twice, and *u* once, and that's all. That's eight *sp*'s, altogether. Now imagine rolling an eight-sided die, four of whose sides are labeled *a*, one labeled *i*, two labeled *o*, and one labeled *u*. Half the time you'll get a side labeled *a*, but half the time you'll get something else. Take that letter, whatever it is (*a*, let's say) and tack it onto your output stream, giving *spa*. Now your most recent digram is *pa*. Once again, consult the statistics telling which characters followed occurrences of *pa* in the input text, and based on these data, make an *n*-sided die and roll it, telling you which new character to append to *spa*—perhaps *s* this time. Now repeat the probabilistic process, this time with digram *as*. Each time, you will get a letter that really occurred in the input text, thus turning a digram into a trigram. And thus, left to right, letter by letter, character by character, fresh new input-imitating output text is generated in this most simple of stochastic ways.

You want some fake state names for the board game you're inventing? Well, take your pick from the following, which were generated by the trigram-frequency method from the database of the fifty real state names (and I could give you hundreds more fake state names at the drop of a hat):

Nebrado, Wessissippi, Oklawaii, New Yornia, Pentahoma, South Delabama, New Jersetts,
Pennectico, Texichussetts, New Hampshington, Michigansas, Oklaware . . .

Now this is recombinant language, in spades!

Using 8-grams to Ape the Great Art of the Bard

Speaking of “recombinant,” back in the 1960s when DNA was quite a novelty, there was a psychologist who trained planarians (a primitive type of aquatic flatworm) on a simple food-finding task, and then ground them up and fed the resultant worm goulash to other planarians. His hope was that the learning that had taken place in the brains of the victimized worms had somehow been stored as sequences in their DNA and, as such, would hopefully survive the chopping-up process, so that the cannibals, after partaking of their worm-goulash banquet, would suddenly find themselves more skilled at finding food. This method of passing on one’s learning to others was at first thought to work and made quite a splash in the press, but later was thoroughly discredited. Oh, well.

With trigrams, new state names
Can be spewed on demand.
The next stop? Why, Shakespeare!
Can his spirit be canned?

Do you remember the sixties?
Some worms were made wise,
Then chopped up and fed
To their relatives in disguise.

The wisdom, it was claimed,
Came through quite intact,
Though the poor DNA
Had been chopped up and hacked.

So let’s chop up old Will
Into 7-grams, or 8—
Will’s wisdom we’ll distill;
Then we’ll regurgitate!

This sounds like a joke—
But take text that’s opaque,
And with trigrams you’ll make
A respectable fake.

When n -grams are used,
 As you jack n up high,
 The facsimile's fakeness
 Gets harder to spy.

Well, Not Quite the Bard ...

The German philosopher Georg Wilhelm Friedrich Hegel produced text that to some folks seems buzzing with wisdom. In my experience, however, most people, at least on their first several hundred readings, find the following Hegel sample fairly unenlightening:

Further, the living substance is being which is in truth subject, or, what is the same, is in truth actual only in so far as it is the movement of positing itself, or is the mediation of its self-othering with itself. This substance is, as subject, pure, simple negativity, and is for this very reason the bifurcation of the simple; it is the doubling which sets up opposition, and then again the negation of this indifferent diversity and of its antithesis. Only this self-restoring sameness, or this reflection in otherness within itself, not an original or immediate unity as such, is the true. It is the process of its own becoming, the circle that presupposes its end as its goal, having its end also as its beginning; and only by being worked out to its end, is it actual.

What happens if you apply the n -gram-frequency method to this text? We can use trigrams, tetragrams, or higher. Here is what came out of a run using pentagrams:

Furtherness with itself, or immediate unity and is in truth substance is the doubling its end also as its self-restoring itself-restoring worked out to its own becoming, the mediation of position of the bifurcation of the bifurcation of this substance is, as subject, pure, simple; it actual.

Note that, even though the program knows absolutely nothing about English words, there is not a single nonword here. On the other hand, the grammar leaves something to be desired. But that gives us good reason to jack up the n -value to, say, $n = 6$.

Further, the movement of position, and only by being which is in truth actual only in so far as its goal, having sameness, or this reflection of its own becoming, the doubling with itself, not an original or immediation in othering which is in truth actual.

Now this time, curiously, we did wind up with two nonwords—"immediation" and "othering"—but they fit in so nicely—they seem so *Hegelian*, one might say—that few readers ever notice that they “do not compute.” But let us now set n equal to 9 and see what happens ...

Further, the living substance is being which is in truth actual only in so far as it is the doubling which is in truth subject, pure, simple negativity, and is for this reflection in otherness within itself, or is the mediation of the simple; it is the true. It is the doubling which is in truth . . .

The grammar is surely getting better, although it is not yet perfect. But grammar aside, what about *meaning*? Well, it sounds more and more like Hegel. Of course, we must beware of that slick phrase “sounds like.” To adapt Bernie Greenberg’s phrase, “It not only should *sound* like Hegel, it should also *speak* like him.” And yet sadly, for us philistines who can’t hear the wisdom resounding in the original lines, these regurgitated lines somehow *do* seem to “speak like Hegel.” And if, to your taste, $n = 9$ isn’t good enough to get credible Hegelian semantics, then jack n up further. Each time you add 1 to n , your output will get a little more realistic, a little more coherent. And please note: this method for simulation of text is far simpler than Emmy, for it is *purely local*, whereas Emmy combines both local and global constraints in a far subtler manner.

Composing in Your Sleep . . . or in Your Grave

Anyone who has composed a fair amount of music knows the exquisite joy of finding their own personal voice. One dreams of composing more and more, but of course time presses and one’s finiteness constrains one’s output. It is therefore natural to wonder what Emmy would do if fed one’s own music as input. And I, given my close connection with Emmy’s progenitor, could request this favor and actually find out. I submitted a diskette to Dave, containing twelve of my piano pieces (I’ve written around forty, all told), and he in turn fed my pieces into Emmy and started her churning. Promptly, out came pseudo-Hofstadter music!

And in the course of my Santa Cruz lecture, Mary Jane performed both a genuine Hofstadter piece and an Emmy/Hofstadter piece (called “Nope”), with me hearing myself aped for the very first time, in real time before the audience’s eyes. It was delightful to listen as my own harmonies were spat back at me in new and unexpected combinations, although I have to admit that sometimes the “logic” of the flow, such as it was, sounded a bit incoherent. But then, had Hegel been listening to someone reading aloud order-9 imitations of his own text, I suspect he too would have found the flow of its logic a little wanting, here and there.

If the blatherings of George Frederick Hegel
 Are captured to a quite Hei deggree
 By 5-grams, shouldn’t George Frederick Handel
 Succumb to the same skulldugg’ree?

But why tackle just big names?
Let's climb down the ladder
To a far easier challenge:
Let's tackle Hofstadter!

Can we suck out the essence
Of Doug from his notes?
Can we psych out Hofstadter
When he hears garbled quotes?

Once his music we've caught,
That's Step 1 of our plan;
The next step's his *lectures*!
We'll make Doug-in-a-Can!

His verse, it's just patterns,
Just rhymes in a box.
There are seldom surprises
That off-knock your socks.

But even the trick rhymes
Have a formula behind;
We know that that's all
That there is to Doug's mind.

The hard part was making him
(The original Doug);
The easy part's faking him
(A canned, Doug-less Doug).

Immortality, ho!—
Thanks to Cope. How I'll rave
When I can compose
In my sleep—or my grave.

There's no one at home,
Yet the music pours out.
The lights have gone dark;
Still, my spirit soars out!

That's Prokofiev's fate.
The poor chappie expired,

Tenth Sonata half-done . . .
EMI finished it. She's hired!

A Prokofiev expert
Said she'd give a première,
So Cope sent it off,
Thinking, "She's quite a dear!"

Not too long had passed
Ere arrived her reply:
"Prokofiev would *hate* this
As much as do I!"

But why the conditional?
Why hedge, using "would"?
If he *sings* while he's dead,
Can't he *hate* just as good?

For this is my claim,
Though it sounds somewhat droll:
Total style-resurrection
Resurrects one's full soul.

Is Language Intrinsically Deeper Than Music?

You may think that that last stanza was tongue-in-cheek, but no—those precise sentiments have been my recurrent theme, and I'm not about to abandon them now. To delve into such matters more seriously, let me discuss more fully in prose what I hinted at in a couple of stanzas above—the idea of a program à la Emmy producing a spate of brand-new Hofstadter lectures or—let's go whole hog—books. (After all, if an opera by Mahler, who never wrote one, is in the planning stages, why not a novel by Hofstadter, who never wrote one?) What would it take? Would n -gram frequencies with a high value of n , applied to all the previously published Hofstadter books, turn the trick?

Well, you know as well as I do that this would fail ludicrously. That kind of technique doesn't deal with *content*, with *ideas*. It just deals with sequences of letters, and a writer does not deal with sequences of letters. And even if an Emmy-like text-imitation program dealt with more global qualities of its input text, the problem is that *new ideas* would not ever enter the scene. Who could have predicted, given my first few books, that I would next write an 800-page book on poetry translation (*Le*

Ton beau de Marot)? (Hofstadter 1997). There's nothing remotely resembling the manipulation and creation of new ideas in Emmy—and yet the crazy, undeniable truth of the matter is that Emmy's music does at least a decent job of creating “new Chopin” and “new Mozart,” and so on. As Dave himself speculated in the journal entry that starts out his first book, “While never as good as the originals, they will be exciting, entertaining, and interesting.”

Or consider “Prokofiev's Tenth Sonata,” as Dave calls it. In the liner notes to his and Emmy's first compact disk (*Bach by Design*), he wrote the following:

This computer-composed Prokofiev *Sonata* was completed in 1989. Its composition was inspired by Prokofiev's own attempt to compose his tenth piano sonata, an attempt thwarted by his death. As such it represents another of the many potential uses of programs such as Emmy (i.e., the completion of unfinished works).

To me this comes close to blasphemy—and yet let me also add the following remark, to counterbalance that reaction. The first movement of this sonata by Emmy starts out with the actual forty-four measures that Prokofiev himself had completed, and then continues with Emmy's own notes. What happens when measures 45, 46, and so on are encountered? Is it like falling off a cliff? Is there a drastic discontinuity? Well, I would put it this way. Imagine you were reading, for the first time, the genuine Hegel paragraph, and imagine furthermore that it had been extended by a high n -value text-imitation program. Would you instantaneously feel something was fishy when you hit the very first word of the computer's text? Of course not. It would take a line or two, possibly even a dozen, before you said to yourself, “What is going on? I'm even *more* lost here than I was at the paragraph's beginning!” In some sense that is how I hear the Emmy/Prokofiev sonata. There is no sudden drop in quality at measure 45—indeed, it is as smooth a transition as one could possibly imagine, and all the way to the movement's end it sounds quite consistent with the way it started out.

I happen not to *like* this piece by Emmy (though Dave Cope adores it!), but then I also happen not to be a fan of the last several Prokofiev piano sonatas (despite loving the first few). So to me, this “tenth sonata,” though it rings fairly true, just doesn't appeal. I would never call it a work of genius, but I would credit it as being “well-crafted and Prokofievian in feel.”

So what is going on? How come Emmy does a fairly passable job at resuscitating composers but couldn't conceivably resuscitate a writer—even a writer of murky and obscure philosophical verbiage? Or am I—a serious author but not a serious composer—being too vain? Are my remarks self-serving? Is my personal verbal spark every bit as susceptible to being captured via algorithmic processing as Prokofiev's musical spark is?

Falling Hard for a Pattern of Blips

For me, analogies always help to shed light on such complex disputes, and the following analogy, which I hope is provocative, came to me one day—indeed, it hit me with great emotional power, I must admit—when I haphazardly picked up a publicity brochure for books about budding new technologies to be put to use in making movies.

I suppose a good deal
Of my EMI perplexity
Can be traced back to issues
Of algorithmic complexity.

To cast all these matters
In a somewhat new light,
Let's turn to attraction,
To chemistry, to "Miss Right."

I recall, and with pain,
A few times in my past
When I fell for some actress
In a romantic film's cast.

The blips on the screen,
We all know, came from her—
A flesh-and-blood person,
Alive, sure as sure.

The image conveyed a full
Human behind the scene,
And *that's* what I fell for—
Not for blips on the screen.

But now let's imagine
A brave new film world
In which love scenes take place,
Both unboy'd and ungirl'd.

And how would this happen?
Quite simple—by CAD—
The faking of objects,
As in many an ad.

One sees things in motion
That in truth never were;
They're simply bit-patterns
Cranked out in a blur.

Of course it's one thing
To make balls bounce about;
Quite another, a person
To believe in, no doubt.

And yet we are marching
Down that very lane;
We're making CAD filmstars.
Is that not a gain?

At this point in my lectures, I generally show the cover of a book called *Synthetic Actors in Computer-Generated 3D Films* by N. Magnenat Thalmann and D. Thalmann (1990), which features a nearly believable but clearly synthesized set of images of a woman in a bathing suit, who is instantly recognizable as Marilyn Monroe, walking across a shiny floor (see figure 2.4); I then read out loud the following blurb for the book:



Figure 2.4

A nearly believable but clearly synthesized set of images of a woman in a bathing suit, who is instantly recognizable as Marilyn Monroe, walking across a shiny floor. (From Thalmann, N. Magnenat and D. Thalmann. 1990. *Synthetic Actors in Computer-Generated 3D Films*. New York: Springer-Verlag.)

Three-dimensional synthetic reincarnations [reader: note this word!] of Marilyn Monroe and Humphrey Bogart were created by the authors of this book for their award-winning feature film “Rendez-vous à Montréal.” The advanced computer animation techniques developed for the film are fully described in this book. They form a technological breakthrough that can be used to produce scenes featuring any celebrity in any situation. This opens new vistas in motion pictures, television, and advertising.

To conjure up Monroe,
Just write code—40 K;
She’ll then dance on your screen,
Blow a kiss, make your day.

Or morph her with Binoche—
Hey, I’d call that a coup!
But since they’re just code,
It’s a no-sweat morpheroo.

So now I can fall
For a screenful of blips
Behind which there’s no one—
Just code caught in chips.

A human in 40 K bytes,
Now that’s cheap—
And yet I might fall
For this “her.” A great leap!

Or *is* it so great
To be gulled by the spiel
Of some code that’s a billion
Times simpler than real?

It shakes me to think
That someday I might fall
For an “actress” who
Never existed at all.

Three Flavors of Pessimism

Yes, what worries me about computer simulations is not the idea that we ourselves might be machines; I have long been convinced of the truth of that. What troubles me

is the notion that things that touch me at my deepest core—pieces of music most of all, which I have always taken as direct soul-to-soul messages—might be effectively produced by mechanisms thousands if not millions of times simpler than the intricate biological machinery that gives rise to a human soul. This prospect, rendered most vivid and perhaps even near-seeming by the development of Emmy, worries me enormously, and in my more gloomy moods, I have articulated three causes for pessimism, listed below:

1. *Chopin* (for example) is a lot shallower than I had ever thought.
2. *Music* is a lot shallower than I had ever thought.
3. The *human soul/mind* is a lot shallower than I had ever thought.

To conclude, let me briefly comment on these. Pertaining to (1), since I have been moved to the core for my entire life by pieces by Chopin, if it turns out that Emmy can churn out piece after piece that “speaks like Chopin” to me, then I would be thereby forced to retrospectively reassess all the meaning that I have been convinced of having detected in Chopin’s music, because I could no longer have faith that it *could only have come from a deep human source*. I would have to accept the fact that Frédéric Chopin might have been merely a tremendously fluent artisan rather than the deeply feeling artist whose heart and soul I’d been sure I knew ever since I was a child. Indeed, I could no longer be sure of *anything* I’d felt about Frédéric Chopin, the human being, from hearing his music. That loss would be an inconceivable source of grief to me.

In a sense, the loss just described would not be worse than the loss incurred by (2), since Chopin has always symbolized the power of music as a whole to me. Nonetheless, I suppose that having to chuck *all* composers out the window is somehow a bit more troubling than having to chuck just *one* of them out.

The loss described in (3), of course, would be the ultimate affront to human dignity. It would be the realization that all of the “computing power” that resides in a human brain’s 100 billion neurons and its roughly ten quadrillion synaptic connections can be bypassed with a handful of state-of-the-art chips, and that all that is needed to produce the most powerful artistic outbursts of all time (and many more of equal power, if not greater) is a nanoscopic fraction thereof—and that it can all be accomplished, thank you very much, by an entity that knows nothing of knowing, seeing, hearing, tasting, living, dying, struggling, suffering, aging, yearning, singing, dancing, fighting, kissing, hoping, fearing, winning, losing, crying, laughing, loving, longing, or caring.

Playing the game of pattern and pattern alone will turn the whole trick—or, as the late and witty mathematician Stanislaw Ulam once said, memorably paraphrasing

Martin Luther (and J. S. Bach), “A Mighty Fortress Is Our Math.” The only difference would be that Dave Cope would proclaim, “A Mighty Fortress Is Our Mac.” And, although Kala Pierson and many others may hail its coming as “truly a thing of beauty,” the day when music is finally and irrevocably reduced to syntactic pattern and pattern alone will be, to my old-fashioned way of looking at things, a very dark day indeed.

If the basis for EMI
Turns out to be true,
Then all my dear notions
Will die; I’ll be blue.

Upon hearing an étude,
I’d no longer conclude
That I sensed a heart’s mood;
’Twas just some “pattern dude”!

Likewise Bach would be shown
To be one “pattern guy,”
Whose secrets are none—
At least not to E.M.I.

Are these two just shallower
Than ever I’d thought,
Their styles simply patterns
In EMI’s net caught?

Or is music itself
Just one big formal game,
So that using brute force
You can ape any name?

Or—worst of my nightmares—
Can a full human “I”
Be stamped on a chip made
By VLSI?

Now don’t get me wrong—
I maintain we’re machines!
But PC’s?! What a slap
In the face to our genes!

Do our millions of genes,
And our billion-celled brains,
Yield nothing but rule-bound
Algorithmic refrains?

I'd like to believe
That for music to spring
From a thing, it must strive,
It must struggle, to sing.

It must search and must seek,
Sometimes win, sometimes fail;
It must fight with the world—
If that's so, I'll not rail.

What I fear is a win by
An emotional sham—
A musical poet with
No sense of "I am."

These issues alarm me
And that's why I spoke—
Not to answer all questions,
But to prod and provoke.

And now, please excuse me
For all of my pranks;
And to Dave (and to EMI!)
I express profound thanks.

Note

1. The epigram was actually stolen from David Cope, who himself had borrowed it from Pablo Picasso.

3

Response to Hofstadter

Doug Hofstadter's highly readable and articulate accounting of how Experiments in Musical Intelligence works, while occasionally overly simplified and understandably incomplete at times, provides an accurate account of the fundamentals of the program's processes. Interestingly, however, Doug's quite vivid and even compelling terminology in describing the program's processes often does not coincide with mine. For example, Doug uses terms like "chop up" and "reassemble," where I much prefer "recombinancy." Where Doug uses terms like "voice-hooking," I prefer to use the more common "voice-leading." His definition of SPEAC, while it reveals Experiments in Musical Intelligence's overall approach, is somewhat limiting, as will be seen in chapter 6. On the other hand, Doug's coinage of the term "templagiarism" describes quite well what I call "unifications."

Some Basics

Doug occasionally admits to an understandable lack of clarity on certain issues. For example, he confesses that he doesn't

... have a clear understanding of how the very complex operation of templagiarism (or, for that matter, the somewhat simpler operation of insertion of signatures) is made to coexist harmoniously with the previously described syntactic and semantic meshing operations, because I can easily imagine them conflicting with each other.

Templagiarism and the placement of signatures occur in parallel with recombinancy rather than sequentially as some (not Doug) misunderstand. In other words, signatures are not actually inserted *after* recombination, but occur as an integral part of the recombination process. Questions about templagiarism, as well as other questions that arise in Doug's presentation, will, hopefully, be resolved in my own more detailed presentation of how Experiments in Musical Intelligence works in chapters 4 through 6.

I find Doug's description of how the program chooses new recombinations of music based on their "shapes" (local connectivity) and their "stuff" (structural considerations) wonderfully poetic and appropriate:

These could be likened to two types of constraints that a jigsaw-puzzle solver naturally exploits when putting together a jigsaw puzzle ...

In one short sentence he has captured the often complex relationships between local and global issues with a visualization that is at once apt and insightful. Unfortunately, this analogy also suggests that there is but one solution, as is the case with jigsaw puzzles. In contrast, Experiments in Musical Intelligence provides many different solutions.

Doug also comments on pattern-matching, particularly in connection with the variables that relate to pattern-matching:

Cope tended to do this adjustment of controllers himself in order to increase the effectiveness of Emmy's search for signatures, but perhaps by now he has managed to automate that aspect of the process.

I have automated this process (see Cope 1992b and 1996, particularly chapter 7) and *Experiments in Musical Intelligence* now routinely composes in this manner.

Doug's view of *Experiments in Musical Intelligence*'s music, as in his description of his first brush with the program's output, often contains somewhat loaded—to me at least—terminology:

It sounded like Emmy was dealing only with the surface level of music—with patterns, not with the deep emotional substrate—and I was pretty sure that little of interest could come of such an architecture.

The words “emotional substrate” sound informational, yet I have only a vague notion of his intended meaning. Likewise,

it [an *Experiments in Musical Intelligence* Chopin work] was unmistakably *Chopin-like* in spirit, and it was *not emotionally empty*.

I am flattered by these comments and yet strangely confused by them as well. Alan Turing (1950) discusses emotion in his seminal article on the Turing test when he quotes from “Professor Jefferson's Lister Oration” for 1949 (the italics are mine):

“Not until a machine can write a sonnet or *compose a concerto* because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that machine equals brain—that is, not only write it but know that it had written it. No mechanism could feel (and not merely artificially signal, an easy contrivance) pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes, be charmed by sex, be angry or depressed when it cannot get what it wants.” (pp. 445–6)

and

According to the most extreme form of this view the only way by which one could be sure that a machine thinks is to *be* the machine and to feel oneself thinking. One could then describe these feelings to the world, but of course no one would be justified in taking any notice. Likewise according to this view the only way to know that a *man* thinks is to be that particular man. (p. 446)

Turing claims, however, that this is “a solipsist point of view” (p. 446). He finally concludes: “I think that most of those who support the argument from consciousness

could be persuaded to abandon it rather than be forced into the solipsist position” (p. 446). Of course, Doug never *actually* argues that Experiments in Musical Intelligence must be conscious to create emotionally satisfying music, just that its music does not deal with music’s “deep emotional substrate.” I have drawn what seems to me to be an implicit connection between emotion and consciousness.

Doug seems to retreat from this point of view somewhat when he writes

I guess I would have to say, “Emmy’s compositions are precisely as deep as her capacity to listen to and understand music.” Which forces us to ask, “Does Emmy listen to music, or understand music, at all?” Well, of course Emmy doesn’t *hear* music, in the sense of having eardrums that vibrate in response to complex waveforms—Emmy’s way of perceiving music involves just the ability to deal with *numbers* that represent pitches and times and so forth.

Doug often seems to equate deep and complex with “good” and “human.” Surely simple things can also be “good” and “human.” At the same time, Doug describes another view that I find compelling, one that involves The Game, which he often includes in his presentations on Experiments in Musical Intelligence. I very much like his rationale for doing so:

I don’t like dishonesty, but perhaps it is best to misinform people about what they are about to hear, in order that they not listen with a preclosed mind.

These lines strike very close to the heart of what I feel occurs when audiences hear the program’s output for the first time, knowing that the music they are hearing was in fact composed by a computer program. Interestingly, I have many times heard a piece of music for the first time, been moved by it, and *not known* who composed it. I typically try to discover the composer’s name in order to place the music in context and, more important, so I can hear more of that composer’s music. However, I have yet to change my mind about the music after discovering who—or, in the case of a computer program, what—composed it. I am more likely to change my mind about the talent of the composer than I am about the quality of the music.

I listen to the Experiments in Musical Intelligence program’s output in much the same way that I listen to any music. I hear the idiosyncrasies of the performance, the mold of the music, and the logic or illogic of the musical anticipations and their consequences. I see no reason to place Experiments in Musical Intelligence’s music outside of the frame of reference that I use for music composed by human beings. On the other hand, I also listen to the sound of wind through a wheat field and the sounds of distant thunder as music. While it is true that I do not listen to these sounds in quite the same way as I do human-composed sounds, they are still music to my ears.

Doug disagrees (see also chapter 16):

I think it is sloppy thinking to equate babbling brooks and birds chirping at twilight with music (at least with traditional tonal music), which is produced deliberately by a human being in order to express or communicate something to other humans.

Doug makes serious assumptions about composers here. Some of the most obvious examples of deliberate attempts at such communication classify as “program music,” which I doubt Doug would include in the list of music he admires. In the tonal music that I have composed in my lifetime, I cannot say that it occurred to me to view myself as deliberately trying to “communicate something to other humans.” I am much more interested in creating well-balanced structures within which I hope to weave inventive musical ideas. I view the output of Experiments in Musical Intelligence in the same light. Doug, however, points out that I am a special case, being the person who created Experiments in Musical Intelligence, and like a proud father, taking great joy in its output. I agree, to the extent that my investment in Experiments in Musical Intelligence has been significant and that I am not dispassionate about its output. At the same time I do not believe that this is the sole reason for my often finding this output musical and moving. I offer the following as possible proof of this.

In the spring of 1998 I joined composers Christopher Dobrian and George Lewis and visual artist Harold Cohen on a performance tour of University of California campuses (which included performances of the Experiments in Musical Intelligence’s Bach and Schubert songs in appendix D). Harold Cohen is best known for his creation of Aaron, his computer drawing and painting program. During Harold’s presentations, which were generally the same from presentation to presentation, as were Chris’s, George’s, and mine, I marveled at his program’s creations. Aaron works quite differently from Experiments in Musical Intelligence in that Aaron does not have a database of previous artworks nor does it create using recombinative techniques. However, Aaron does produce artworks that resemble those created by humans, artworks that sometimes sell for \$5000 apiece. Figure 3.1 shows a black-and-white version of one of Aaron’s color paintings which I find particularly evocative—and one created especially for this book. The figure in this picture reminds me of someone I know and this figure’s expression reveals many interesting contradictions. In short, my relationship to Aaron’s art is precisely the same as if Harold or any other human painter had painted it. My “privileged” relationship to Experiments in Musical Intelligence does not equally connect me with Aaron so that I can say, without reservation, that for me at least, computer-created art in any form, while it must pass my own personal muster of quality, falls clearly within the same




Figure 3.1

A black-and-white version of a color painting by Harold Cohen's Aaron program.

frame of reference as does human-created art (see further discussion of these points in chapters 17 and 18).

Doug's comments on computer-generated text involves a particularly opaque quote of Hegel's which he describes thusly:

... most people, at least on their first several hundred readings, find the following Hegel sample fairly unenlightening.

He then proceeds to use the quote as a template for n -gram-frequency generations. After a series of attempts he writes:

Well, it sounds more and more like Hegel. Of course, we must beware of that slick phrase "sounds like." To adapt Bernie Greenberg's phrase, "It not only should *sound* like Hegel, it should also *speak* like him." And yet sadly, for us philistines who can't hear the wisdom re-sounding in the original lines, these regurgitated lines somehow *do* seem to "speak like Hegel."

This is *very* revealing. Doug seems to intimate that if it is complex enough, then it will actually sound as if it is deep, even if it isn't deep at all. As I mentioned earlier, Doug often seems to relate deep and complex to "good." I do not. Some of the most moving and profound experiences of my life have been some of the simplest and most naive.

The Source of Frustration

In all of Doug's prose and doggerel (his term), he seems fixed on three possibilities which he poses, *each* of which has to do with deepness and shallowness:

1. *Chopin* (for example) is a lot shallower than I had ever thought.
2. *Music* is a lot shallower than I had ever thought.
3. The *human soul/mind* is a lot shallower than I had ever thought.

Doug's frustration peaks when he is confronted by a work of Experiments in Musical Intelligence that seems as if it had been composed by a once-alive composer he respects. I am again, of course, delighted that the program's output seems to be of a quality that would produce such a reaction. However, it seems to me that depth of music, computer programs, or people are not the true source of Doug's frustration. The problem arises, it seems to me, from identifying the wrong source of at least part of the depth.

For example, a woman once remarked to me that she felt certain that a computer program could compose music as beautiful as that composed by humans. However, she also felt that human-composed music contains hidden messages that intrigued her to decipher. The combination of beauty and revelation from such decoding in her mind made human-composed music greater than machine-composed music. She allowed that computer-composed music could also contain hidden messages, but stated that she was not interested in decoding these messages because they were unintentional. I replied that both types of messages were of her own making and that music from both sources had intention: *her* intention. Stravinsky remarked:

... music is, by its very nature, powerless to *express* anything at all, whether a feeling, an attitude of mind, a psychological mood, a phenomenon of nature ... if, as is nearly always the case, music appears to express something, this is only an illusion and not a reality. It is simply an additional attribute which, by tacit and inveterate agreement, we have lent it, thrust upon it, as a label, a convention—in short, an aspect unconsciously or by force of habit, we have come to confuse with its essential being. (Stravinsky 1975, pp. 53–4)

John Cage echoes these thoughts when he observes that

Most people think that when they hear a piece of music, they're not doing anything but that something is being done to them. Now this is not true, and we must arrange our music, we must arrange our art, we must arrange everything, I believe, so that people realize that they themselves are doing it, and not that something is being done to them. (Nyman 1974, p. 21)

Listeners obviously play a significant role in the musical experience. More than that, however, listeners play a *primary* role in this experience. We love or hate music not because the music itself is lovable or hatable but because we, through a complex and probably impossible-to-define set of aesthetic processes, decide to love it or hate it.

However, I don't mean to imply that music itself does not contain meaning. Consciously *and* subconsciously, composers invest in music aspects of everything they

have heard and understand about music. After all, composers compose *recombinantly*. I use this term deliberately, since I believe Experiments in Musical Intelligence uses processes of recombination similar to those that human composers use to compose. I believe as well that these same processes create the meaning we hear in music (more on this later).

Doug once confided in me that he had dreamed that I had confessed to composing Experiments in Musical Intelligence's entire output myself. This dream apparently produced a great feeling of relief in him. Since Experiments in Musical Intelligence has now composed over six thousand works, I doubt that his dream has recurred. This dream, along with Doug's three possibilities of deepness in music, suggests to me that he feels the quality of Experiments in Musical Intelligence's works somehow *cheapens* the works of his favored composers. To me, however, the flaws in the program's output help me *appreciate* the human-composed works I love more than ever. Remember that no Experiments in Musical Intelligence music would exist without the originals in the database—a fact that in itself should dilute any imaginings that these works supersede their predecessors in quality.

Doug's final assertion—"The *human soul/mind* is a lot shallower than I had ever thought"—is wrong. I think that we are all a lot *deeper* than even he imagines us to be. And I, for one, am not at all convinced that this is necessarily such a good thing.

Doug's idea of musical depth (see also chapter 16) seems intertwined with our Western tradition of ascribing "greatness" to some composers while others are of lesser quality:

... to claim that the active involvement of our recipient brains transfers all credit for greatness and depth from the creator to the recipient is nonsense. Such a viewpoint would imply that we can find greatness in anything at all, simply by using our receiving brains' power.

This quote raises some very important issues, one of which, the role of the listener, I address extensively in chapter 18. Here I wish to address the notion of "greatness." I do not believe that any work of art is intrinsically better than any other work of art. Greatness is something we impose on certain favored composers or works, not something that resides implicitly within them. Greatness arises and passes only by the whims of social, cultural, and historical fashion, not because works of art are *actually* great. Greatness is an opinion, not a fact. Unfortunately, in our rush to lionize certain composers we forget the extraordinary achievements of many less fortunate but nonetheless highly gifted lesser-known composers. This seems so self-evident to me that I find it difficult to articulate, and fear that by saying it I will insult my reader's intelligence. It is because I find many of my professional colleagues in disagreement that I risk raising the point here.

A friend of mine often asks the question: “Why are squirrels so damned cute and rats so damned ugly? They are both rodents. They both carry disease. Is it just because the squirrel’s tail is so damned cute and the rat’s tail so ugly?” The answer, he claims, is simple: “The squirrel’s tail *is* damned cute.” I do not yet know of anyone who favors the rat in such a comparison, nor anyone who, when asked to truly evaluate their answer to this question, actually believes that there is something intrinsically better about squirrels than rats.

Soul

Doug often grieves of his being “moved by 20,000 lines of code” (his way of describing Experiments in Musical Intelligence). By abstracting the program in this way, Doug focuses our attention on the inspiration-less, imagination-less, and soul-less form that he perceives the program to be. What this image (not Doug) fails to take into account is the thousands of hours of human labor that went into the program’s creation. The fact is, any words or musical notes spread out on a sheet of paper look as cold and lifeless as do 20,000 lines of code. It is only our imaginations that breathe life into those markings. My 20,000 lines of code are, after all, simply *instructions*. Hemingway’s instructions involve words to read. My instructions relate to how I want my computer to act. While my instructions may seem once removed from the actual creation of a new work of art, they are just as real and just as filled with my hopes for success as are Hemingway’s instructions.

As Doug notes (see his comments in chapter 16), I dedicated a mazurka in the style of Chopin by Experiments in Musical Intelligence to him. My hope was twofold: to thank him for his continued efforts to make this body of work visible and audible in his concert-lectures and to personalize for him what was obviously an impersonal process. This mazurka appears, along with its dedication, in appendix D. My emphasis over the years with Experiments in Musical Intelligence has been to focus on its output—the music—rather than the source, the computer program. In so doing, I hope to underscore what I feel, as a composer, is the program’s most important contribution: interesting and, to me at least, occasionally deeply moving music.

The essence of Doug’s frustration appears in his own words, and again seems intimately connected with his notions about depth:

What troubles me is the notion that things that touch me at my deepest core—pieces of music most of all, which I have always taken as direct soul-to-soul messages—might be effectively produced by mechanisms thousands if not millions of times simpler than the intricate biological machinery that gives rise to a human soul.

The question of “soul” surfaces almost every time someone hears and then responds to a work by Experiments in Musical Intelligence. To begin with, I always try to start with a good definition of *soul*. *Webster’s Dictionary* (1991, p. 1278) defines *soul* as “The principle of life, feeling, thought, and action in humans, regarded as a distinct entity separate from the body.” There are three particular aspects of this definition that clearly exclude computers and hence computer creations from having soul. First, computers do not have life, by any current biological definition of that word. Second, computers are not human. Third, by inference at least, computers do not contain anything separate from their bodies, nor do I imagine that they will have any such separate “thing” in the near future. Therefore, computers do not have soul by definition. Nor can they ever have soul by such an exclusionary definition.

The bigger question for me, however, is not whether Experiments in Musical Intelligence compositions have soul but whether *human* compositions have soul (note that I deliberately avoid the question of whether humans themselves have soul). With such a vague definition of soul, I cannot see anyone arguing convincingly that the notes gathered on a page of music contain in them a “principle of life, feeling, thought, and action.” No, the soul we perceive when we hear a deeply moving musical work, if “soul” is even the right word, is *our own soul*.

I believe the fact that we cannot, in general, recognize the difference between a machine-created and human-created work of art means that the program uses processes that in some ways mirror those used by human composers. In fact, I make the case for this in Cope 1996. Doug argues (see also chapter 16):

I would take issue ... that a composer starts out spending a good deal of time devising a grammar, and then, that having been done, just turns into a drone who spins off piece after piece using the rules of the grammar.

Doug is a master of language and here he wields “drone” and “spins off” as natural consequences of what he describes. I would argue that Experiments in Musical Intelligence is *not* a drone and that it *composes* rather than spins off new works.

However, Doug’s commentary often proves quite insightful, pointing out aspects of the program that I rarely if ever mention, but should. First, he comments on how the program is not singular but plural:

It is indeed true that Emmy is evolving—it is a moving target.

This is a very important point. Many lovers and critics of Experiments in Musical Intelligence works seem to feel that there exists one single stable version of the program, when in fact there have been many (perhaps dozens of) versions of the program. Were it not for the fact that I keep rigorous documentation, it would be

difficult for me to account for the differences between various Experiments in Musical Intelligence outputs.

Doug also notes (in reference to the creation of an Experiments in Musical Intelligence work in his own style in chapter 16) that

if you have such a sparse database on which to draw, your output is going to reek of its sources in a conspicuous manner. If Dave had been able to use all twelve of the pieces I sent him (or better yet, all forty of the pieces I've composed), then of course the mixture would have been far, far subtler.

I agree. A composer I know went to one of Doug's presentations on Experiments in Musical Intelligence. During the question period following the presentation he commented to Doug: "You should spend your time speaking about the possibilities of computers being used to compose *contemporary* music rather than all of these style imitations." For me, this anecdote accentuates the controversy created by Experiments in Musical Intelligence. There doesn't seem to be a single group of people that the program doesn't annoy in some way.

Interestingly, Doug's experiences with audiences is inversely proportional to my own experience.

Hardly anyone seemed upset at Cope's coup in the modeling of artistic creativity; hardly anyone seemed threatened or worried at all.

Doug's presentations, at least the ones that I have attended, range from delightfully humorous to almost painfully emotional. My presentations, on the other hand, cover the nuts and bolts of how the program works, usually without any references to emotion or to the program's philosophic implications. I suppose that audiences tend to pale when faced with his angst while they anger at my apparent detachment.

When Doug says (see chapter 16) that

Emmy's power of course comes, in *some* sense, from borrowing, for by definition and by intention, that is all that Emmy is—a borrower.

I must remind him of the quote he uses at the beginning of his commentary:

Good artists borrow; great artists steal.

While on this topic, I note that Doug's endnote is mistaken:

The epigram was actually stolen from David Cope, who himself had borrowed it from Pablo Picasso.

I did not *borrow* this line from Picasso, I *stole* it from Picasso, who actually *borrowed* it from Stravinsky.