

Virtual Music

Computer Synthesis of Musical Style

David Cope

With commentary by Douglas Hofstadter

*And with perspectives and analysis by Eleanor Selfridge-Field, Bernard Greenberg,
Steve Larson, Jonathan Berger, and Daniel Dennett*

The MIT Press
Cambridge, Massachusetts
London, England

15 Collision Detection, Muselot, and Scribble: Some Reflections on Creativity

Daniel Dennett

In this chapter I want to stress the continuity between all sorts of creativity. Life on earth has been generated over billions of years in a single branching tree—the tree of life—by one algorithmic process or another. Figure 15.1 shows a recent picture of the tree of life (Morell 1997, p. 701). Note that a relatively recent branching on the eukaryotic limb, which bears all multicellular organisms, divides the plants from the animals. Among the plants are apple trees, which make apples, and among the animals are spiders, which make webs, and beavers, which make dams. Apple trees and apples, spiders and webs, beavers and dams—these are all fruits of the tree of life, directly or indirectly.

Also among the animals are three magnificent composition makers: J. S. Bach, David Cope, and Bernard Greenberg (who performed a number of his works patterned after Bach at the conference). They and their compositions are also among the fruits of the tree of life, and one of these composition makers, David Cope, has also created a composition maker, *Experiments in Musical Intelligence*. Its products are doubly indirect, being products of a product of a product of the tree of life, but, I want to argue, its means of production are simply special cases of the very same processes that created both the compositions by Bach and Greenberg, the apples and spider webs, and the organisms that made them.

According to the poet Paul Valéry,

It takes two to invent anything. The one makes up combinations; the other one chooses, recognizes what he wishes and what is important to him in the mass of the things which the former has imparted to him. What we call genius is much less the work of the first one than the readiness of the second one to grasp the value of what has been laid before him and to choose it. (quoted in Hadamard 1949, p. 30)

This making up of combinations is also known as generating diversity, and the choosing is also known, of course, as selection—as in natural selection. Valéry was right. All invention, all creation, proceeds by trial and error of one sort or another, and all such processes are what are known as generate-and-test algorithms (Dennett 1975). Since the processes are algorithmic, we ought to be able to take them apart and see how they work. How many layers of generate-and-test does it take to create an apple tree, or a Bach, or the *St. Matthew Passion*, or *Experiments in Musical Intelligence*, or one of that program's inventions?

Many find this vision of creativity deeply unsettling—some would add that it is not just unsettling, it is crass, shallow, philistine, despicable, or even obscene. I am always

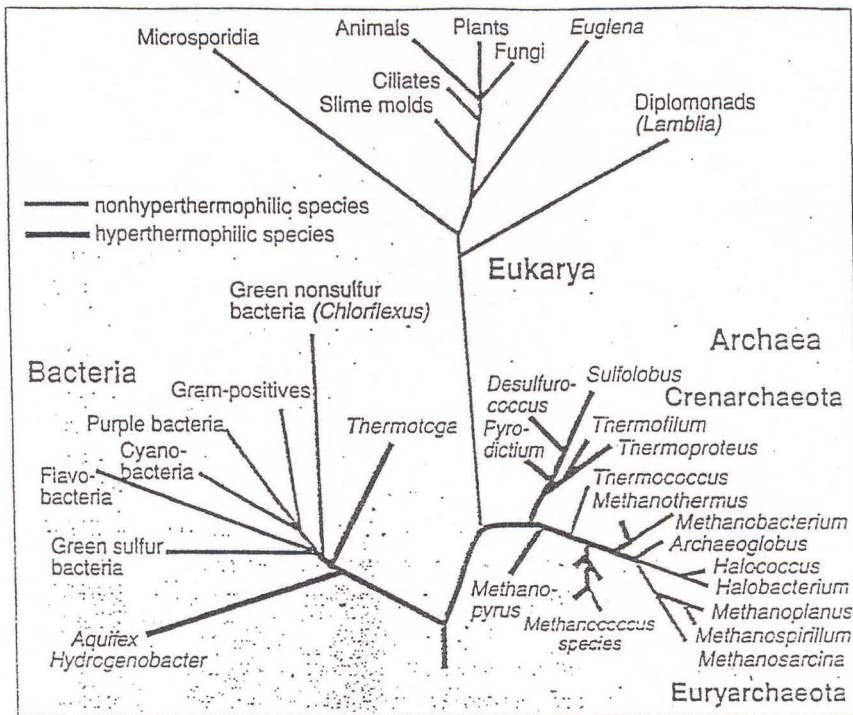


Figure 15.1

A recent picture of the Tree of Life (Morell 1997, p. 701).

more than a little quizzical about these emotional reactions, since they seem to me to betray a certain parochialism. It is apparently *not* crass, philistine, obscene ... to declare that all the first-order products of the tree of life—the birds and the bees and the spiders and the beavers—are designed and created by such algorithmic processes, but outrageous to consider the hypothesis that creations of human genius might themselves be products of such algorithmic processes. Is an ode to a nightingale really that much more marvelous than a nightingale itself? I would certainly be prouder of having created the latter! And of course no human being, no matter how great a genius, does all the creative work that goes into a work of art.

How long did it take Johann Sebastian Bach to create the *St. Matthew Passion*? An early version was performed in 1727 or 1729, but the version we listen to today dates from ten years later, and incorporates many revisions. How long did it take to create Johann Sebastian Bach? He had the benefit of forty-two years of living when the first version was heard, and more

than half a century when the later version was completed. How long did it take to create the Christianity without which the *St. Matthew Passion* would have been literally inconceivable by Bach or anyone else? Roughly two millennia. How long did it take to create the social and cultural context in which Christianity could be born? Somewhere between a hundred millennia and three million years—depending on when we decide to date the birth of human culture. And how long did it take to create *Homo sapiens*? Between three and four billion years, roughly the same length of time it took to create daisies and snail darters, blue whales and spotted owls. Billions of years of *irreplaceable* design work. (Dennett 1995, p. 511)

The contribution of Christianity to Bach's work is, in one sense obvious, but it is made particularly vivid by Bernard Greenberg (see chapter 12), who has analyzed his own reliance on the memes of Christianity to enrich his own Bach-inspired compositions. For those interested in theories of cultural evolution, Greenberg provides a particularly clear-cut example of a type of cultural vector (and locus of combination) that has been overlooked by many theorists: he is a meme vector, and meme exploiter, without being a believer. He harbors an elaborate virtual machine, the well-designed, field-tested, debugged culture machine of Christianity, and he has gone to considerable lengths to acquire this machinery. It enables him to do things that he could not do if he hadn't incorporated that virtual machine into his architecture, but he does not believe its doctrines. You don't have to believe in a system of memes to transmit them, or to benefit from them. (This has been typically overlooked by Darwinian theorists of culture. For some details, see Dennett, forthcoming.) And one of the remarkable contributions of Greenberg's presentation to the conference is his willingness to let us go backstage and see how the process works—for him, and by extrapolation, for others. Unlike Greenberg, most artists conceal their creative processes from us as if they were magicians. Picasso famously said, "Je ne cherche pas. Je trouve," but if he was sincere (which I doubt), this only means that he had less access to, and less insight into, the search procedures (or generation procedures) that actually went into his own formidable artistic production.

Here is a recipe, then, for making the *St. Matthew Passion*. First, make a Bach, and educate it, installing all the best products of the contemporary culture. Then sit back. Pretty soon it will be punching out cantatas one a week, for years. From there it is a large but not miraculous step to a more ambitious work like the *St. Matthew Passion*. J. S. Bach was prolific, but not as prolific as the Experiments in Musical Intelligence program—but then Bach was running on a much slower architecture, using old-fashioned technology. What an algorithm can do in principle, and what it can do in jig time (or just in real time in real life) is entirely a matter of the architecture on which it is run. When Danny Hillis, the founder of Thinking Machines, Inc., wanted to show off the huge advance in computing power of his Connection




Figure 15.2

Evolved virtual creatures (Sims 1994a,b), an artificial evolution program that produced strikingly biological-like phenomena.

Machines, he commissioned Karl Sims to create some demos. One of the best was his Evolved Virtual Creatures (Sims 1994a,b), an artificial evolution program that produced strikingly biological-like phenomena in only a few hundred generations (see figure 15.2). Like *Experiments in Musical Intelligence*, this program is itself the work of a brilliant creator, but also like *Experiments in Musical Intelligence*, its creations include many features undreamt of—even antecedently unimaginable—by its creator. But for all its brilliance—and my admiration for Sims’s work is unstinting—the program does have some striking limitations. By drawing attention to them, I want to suggest a ground for residual misgivings about *Experiments in Musical Intelligence*, which has—to date—the same limitations.

What’s missing in Sims’s world of evolving virtual creatures, is, in a word, *concreteness*. For instance, although the bodies of these virtual creatures are exposed to selective pressure in the virtual world they inhabit, their genomes are offstage, and hence not subject to any selective pressure. That means that their genomes could not grow longer or shorter under any realistic circumstances. While there are a Vast (Dennett 1995) number of possible virtual creatures in Sims’s world, the search space

is open-ended in a strictly limited number of dimensions. When Sims wanted to see if his creatures could evolve phototropism (heading for the light, the way moths close in on lamps and candle flames), he had to reach in, godlike, and add the genes for a photosensor element, increasing the dimensionality of his simulation. Another evolutionary simulation, John Holland's (1995) ECHO, does put the genome itself in the virtual world; organisms in ECHO have to acquire the necessary raw materials to make offspring, including the offspring's genome, before they can reproduce. This dramatically widens the scope of evolutionary possibilities, but still falls short of the open-endedness that only a fully concrete phenomenon can enjoy.

Consider the difference between virtual worlds and real worlds. If you want to make a real hotel, you will have to put a lot of time, energy, and materials into arranging matters so that the people in adjacent rooms can't overhear each other. When you go to make a virtual hotel, you get that insulation for free. In a virtual hotel, if you want the people in adjacent rooms to be able to overhear each other, you have to add that capacity. You have to add noninsulation. You also have to add shadows, aromas, vibration, dirt, footprints, and wear and tear. All these nonfunctional features come for free in the real, concrete world. The generic term for what must be added to virtual worlds to make them more realistic is *collision detection*. If you have ever played around with the task of making a computer video game, you soon realize that putting shapes in motion on the screen is not enough. Shapes will pass right through each other without any effect unless you build collision detection into the update loop.

In his book *Le Ton Beau de Marot*, Hofstadter (1997) draws attention to the role of what he calls *spontaneous intrusions* into a creative process. In the real world, almost everything that happens leaves a wake, makes shadows, has an aroma, makes noise, and this provides a bounty of opportunities for spontaneous intrusions. It is also precisely what is in short supply in a virtual world. Indeed, one of the chief beauties of virtual worlds from the point of view of computer modelers is that quietness: nothing happens except what you provide for, one way or another. This permits you to start with a clean slate and add features to your model one at a time, seeing what the minimal model is that will produce the sought-for effects.

Sims's Evolved Virtual Creatures is a spectacular example of getting a lot from a *relatively* simple model, but it also serves to show that when you're modeling creativity, there should be junk lying around that your creative processes can bump into, noises that your creative processes can't help overhearing. The spontaneous intrusion of that little noise from the next room may tweak what those processes are doing in a way that is serendipitous, or in a way that is destructive, but either way, this opens up

new possibilities. The exploitation of accidents is the key to creativity, whether what is being made is a new genome, a new behavior, or a new melody.

Let me clarify what I'm *not* saying. The problem with Sims's evolved creatures is not that they are not made of carbon, or that they contain no protein or hemoglobin. The problem is that they are virtual. And by being virtual, they live in a world many orders of magnitude simpler than the world of biological evolution. I think exactly the same thing is true of Experiments in Musical Intelligence. Wonderful as it is, it is orders of magnitude simpler than the world of human musical composition. What is delightful about both cases is the discovery of just how much you can get from something so clean, so noise-free, so abstract.

We can imagine improving Experiments in Musical Intelligence, or Karl Sims's work, or any other such project in artificial life or artificial creativity, by adding more and more and more junk, more and more opportunities for collisions, into the world. But consider how counterintuitive such advice would appear:

No matter what you're modeling, make sure that every phenomenon, every subroutine, everything that happens in that world makes extraneous noises, leaves a wake, broadcasts a variety of nonfunctional effects through the world.

Why? What is all this noise for? It's not for anything; it's just there so that every other process has that noise as a potential source of signal, as an *objet trouvé* that it *might* turn, by the alchemy of the creative algorithm, into function, into art, into meaning. Every increment of design in the universe begins with a moment of serendipity, the undesigned intersection of two trajectories that yield something that turns out, retrospectively, to be more than a mere collision. But to the extent that computer modelers follow this advice, they squander the efficiency that makes computers such great tools. So there is a sort of homeostasis here. We can see that, not for any mysterious reason, computer modeling of creativity confronts diminishing returns. In order to get closer and closer to the creativity of a human composer, your model has to become ever more concrete; it has to model more and more of the incidental collisions that impinge on an embodied composer.

In other words, I want to follow Greenberg and Hofstadter in insisting that the shortcomings—such as they are—of Experiments in Musical Intelligence have nothing to do with carbon vs. silicon, biology vs. engineering, humanity vs. robotity. Concentrating on these familiar—indeed threadbare—candidates for crucial disqualifiers is succumbing to misdirection. What, though, about another all-too-familiar theme: what about consciousness? Experiments in Musical Intelligence is not conscious—at least not in any way that significantly models human consciousness. Isn't the consciousness of the human artist a crucial ingredient? One might think that there could

not possibly be any *meaning* in the product of an unconscious process. Interestingly, some of the most creative human beings of all time have already thrown cold water on this hunch. (I have discussed these cases before, in Dennett 1975, from which the following paragraph—with minor revisions—is taken.)

When the great 19th century mathematician, Henri Poincaré, reflected on this topic, he saw only two alternatives and they were both disheartening to him. The unconscious self that generates the candidates “is capable of discernment; it has tact, delicacy; it knows how to choose, to divine. What do I say? It knows better how to divine than the conscious self since it succeeds where that has failed. In a word, is not the subliminal self superior to the conscious self? I confess that, for my part, I should hate to accept this” (Koestler 1964, p. 164). The other extreme, of course, is that the generator is just an automaton, an ultimately absurd, blind trier of all possibilities. And that’s no more of a homunculus with whom to identify oneself. So you don’t want to be the generator. As Mozart is reputed to have said (in an oft-quoted but possibly spurious passage): “Whence and how do they come? I don’t know and I have nothing to do with it.” (p. 86)

One of my favorite contemporary creators, the novelist and essayist Nicholson Baker (1996), has this to say on the way we go about changing our minds:

Our opinions, gently nudged by circumstance, revise themselves under cover of inattention. We tell them, in a steady voice, No, I’m not interested in a change at present. But there is no stopping opinions. They don’t care about whether we want to hold them or not, they do what they have to do. (p. 4)

What these reflections have in common is a vision that many find deeply repugnant, even alarming: the self disappears, or at least shrinks, as Thomas Nagel once said, to a dimensionless point (Nagel 1979, p. 35). The active, responsible, farseeing, genuinely creative self is replaced by willful opinions or unconscious automatic selves that go right on generating and testing *without authorial supervision*. That the generation might be unconscious is recognized by Poincaré and Mozart; Baker adds that even the test is typically something that can happen outside of our control. Looked at from so close up, I seem to vanish. As Doug Hofstadter asks, in the last of his list of worries, “Am I not as deep as I thought I was?” But once asked, the question’s answer is not far to seek. Indeed, another of my favorite novelists, Peter DeVries, has a character express just the right sentiment for this impasse: “Superficially he’s deep, but deep down he’s shallow.”

Of course. How could it be otherwise? Hofstadter recognizes that although his self has many, many layers, it does not have an infinite number of layers, so when our analysis bottoms out, we find nothing but shallow, mechanical processes built on shallow, mechanical processes, built on shallow, mechanical processes. Completing the analysis, even in sketchy, speculative fashion, as we are doing here, should be an

occasion for joy, not despair. Is it not wonderful that a cascade of algorithmic processes should look really deep from a superficial point of view? What else could depth be, in the end? Hofstadter understands this, of course. His disappointment is rather the reaction of somebody who would have guessed that a musical mind had at least, say, seventeen layers, instead of “just” thirteen.

Here, for once, I think I disagree somewhat with Hofstadter. It seems to me that he is hoping to sustain a view of meaning *in music* that is one (or two or three) layers too deep. We can frame the issue by comparing music to speech, and considering a fascinating variety of ersatz speech invented—or at least named—by Dario Fo, the Italian poet, dramatist, and comedian (and 1997 Nobel laureate in literature). Among his comedic tours de force are his exercises in what he calls *gramelot*. The art of *gramelot* is the art of *seeming* to speak a foreign language that one is in fact not speaking. In *Mistero Buffo*, his *gramelot* Elizabethan English judge sounds just like genuine Shakespearean English to Fo’s non-Anglophone Italian audiences—but sounds marvelously strange and alien to us Anglophones. *Gamelot* has been invented time and again by comedians. Danny Kaye and Mel Brooks have done hilarious phony French and German, and then there is the Swedish chef on the Muppet Show, and the exquisite mouthings of the proprietor of the Café Boeuf on Prairie Home Companion.

What, then, is Hofstadter’s worry? It is that Experiments in Musical Intelligence is just producing musical *gramelot*, what we might call *muselot*—and that that is all there is to music! Music just *sounds as if* it really means something! The reasoning that leads Hofstadter to this quandary is quite straightforward.

Premise 1: Experiments in Musical Intelligence’s music doesn’t mean anything, however wonderful it is.

Premise 2: Either there is a *fundamental* difference between the Experiments in Musical Intelligence program’s music and Bach’s (and Schubert’s, and Puccini’s . . .) music or there isn’t.

Premise 3: There isn’t.

(For Hofstadter, as for me, this premise follows from our conviction that strong AI is possible in principle, however impossible it may turn out to be in grubby, economic fact.)

Conclusion: Human music doesn’t mean anything, however wonderful it is.

How could it possibly be that music means nothing, if untold generations of musicians and composers have thought that it does have meaning? This rhetorical question is easy to answer. Consider what linguists sometimes call *scribble*, the mixture of babbling and meaningful words that young children enthusiastically emit for a brief period in their language acquisition process. As has often been pointed out, there is no reason to suppose that these infants have any idea that their production is

falling short *in any way* from what they hear adults saying. After all, from their point of view, what adults are speaking is a lot of gibberish interspersed on occasion by a familiar, meaningful word. In the case of scribble, infants soon learn that even the gibberish part—the impressive gramelot—has meaning, but we can imagine (I think) a quite stable linguistic practice that never went beyond scribble. I harbor the suspicion that the fashionable works of recent French philosophy and literary studies are just such productions, and Hofstadter's examples of randomly generated Heidegger and Hegel indirectly support my dire conjecture.¹ But whether or not a *linguistic* institution of partial meanings embedded in delicious-sounding nonsense could survive, I see no reason why music could not consist of just such a matrix of occasionally (and obviously) meaningful elements set in structures that had no semantic interpretation at all.

The meaning that we know music can have is hardly negligible. Greenberg responds directly to Hofstadter in his discussion of Bach's use of counterpoint as a building system, a structure for incorporating meaning in music. But I think the elements of meaning that he shows to be incorporated in Bach's work still fall short of providing a standard of musical meaning that would permit us to draw a distinction between music and muselot, analogous to the difference between meaningful speech and gramelot. Even if we grant—what is surely true—that there is wide variation in the competence of human listeners to discriminate, abstract, and appreciate the music that they hear, this variation is not radical enough to parallel the imagined race of people who are speaking scribble and not realizing that other people are actually talking. The background fear that there might be a musically elite class that actually performed and appreciated music, while the rest of us were just content with muselot, is an empty fear. It couldn't be true. If it were true, there would be differences in the things some of us could do with music that would be as manifest to us as the differences between a good cookbook and a cookbook composed of randomized recipes. Music can "speak to us" but it can't give us directions, or explain phenomena, or codify laws, or tell stories (without the help of lots of verbal signposts, titles, and labels). Music is deep, music is wonderful, but not *that* wonderful.

Note

1. Doug Hofstadter coined the word "templagiarism" at the Stanford conference; it inspires me to attempt a coining of my own: *eumerdification*. John Searle once told me about a conversation he had with the late Michel Foucault: "Michel, you're so clear in conversation; why is your written work so obscure?" To which Foucault replied: "That's because, in order to be taken seriously by French philosophers, 25 percent of what you write has to be impenetrable nonsense." So, according to Searle, Foucault claimed that he deliberately added 25 percent *eumerdification*, so he would be taken seriously in France.