

Chapter 17

The Future of Music

Technology has already had a significant effect on the business of creating and distributing music. But the development of a complete theory of music will have an even more dramatic impact.

In the future, to get new music, you won't have to wait for someone to write it—you'll just push a button on your computer. This will completely change the economics of the music industry—musical composition and song-writing will cease to be an activity anyone can get paid for, and there will be a premium on artists who can learn and play newly composed music on the fly. Or perhaps we won't bother to wait for human performers to keep up with the machine, and we'll just tell our computers to compose the music and then perform it for us as well.

A musical composition algorithm based on a scientific understanding of music could make music into the newest drug—one that cannot be banned.

17.1 Music as a Commercial Enterprise

Music, like many other aspects of human culture, has been transformed by Western capitalism into a giant business enterprise. The informational nature of music means that small numbers of composers and performers can potentially supply the musical needs of hundreds of millions of music-listening consumers. Enormous resources can be put into the production of music. A record company might spend hundreds of thousands of dollars (US) to produce an album for a band or performer that is expected to sell well.

The corollary of this few-to-many relationship is that consumers can expect to get the best possible music, and only the best, as long as their personal taste is not too distinct from everyone else's.

Cynical commentators will say that the capitalistic nature of the modern Western music industry has some negative effects, that record companies become concentrations of marketing power with an inertia of their own that can resist innovation, that the music industry is distorted by the demographics of who spends on albums and who doesn't, and that the industry creates talentless "manufactured" bands and artists according to a marketing formula.

Despite these criticisms, I am prepared to believe that a basic consequence of the economics of the modern Western music industry and its technologies of production and distribution is that the average modern Western music consumer is exposed to music that contains a much more intense musicality than what would have been heard by a person living in a small tribe thousands of years ago. In the prehistoric tribal environment, "production" would have consisted of people singing and playing hand-made instruments in a small village environment, and "distribution" would have consisted of live performances, and perhaps some transference of musical skills and repertoire from teacher to pupil by means of direct teaching.

The specific features of the modern music industry that enable the development and distribution of high quality music include the following:

- Electronic and computer systems for recording, playing, composing and altering musical sounds.
- Consumers with a high level of disposable income, some of which is available to fund a pool of composers and musicians, from which the *crème-de-la-crème* of composition and performance can be selected.
- Efficient mechanisms of distribution, including radio, printing of sheet music and production of recorded music, which allow almost everyone in society, including those working in the music business, to be exposed to a large body of existing work.
- Enormous potential rewards (for the composers and performers), which encourage the development of ever-improving music for the listening audience. This can sometimes be counterproductive, as too much money and luxury allows successful musicians to either kill themselves from overindulgence, or permanently retire from full-time work. But these effects of excessive wealth are more than offset by the benefits of having the freedom to work full-time (if an artist can make enough money from their music to support themselves), and by the relentless human tendency to create and to achieve (which continues even when one has made enough money to retire on).

The most recent development in this mix is the Internet. The Internet has the potential to put a damper on some aspects of the music industry.

Hundreds of millions of people have ready access to computers connected to a global public network. The invention of efficient lossy compression algorithms—such as **MPEG Audio Layer 3** (more commonly known as **MP3**)—has enabled music fans to quickly and easily (and illegally) transfer near-perfect copies of commercially produced music recordings between themselves, denying performers, composers and producers income for their work.

But it is not all bad news:

- The Internet lowers the bar for distribution and advertising costs. Internet technologies such as email, blogs, RSS, search engines and “social software” (and probably a few more technologies yet to be invented), can streamline word-of-mouth communications so that a new performer becomes known in the shortest possible time to the largest possible audience, with near-zero advertising cost. The problem then remains of how to collect any money after you have become famous. Artists may have to rely to a larger extent on income derived from live performances. Or society as a whole may have to determine a means of paying content producers that does not depend on the ability to enforce copy protection.
- The Internet can help to lower the initial production costs of making music, by making it easy for people creating music to share ideas, content and software with each other. Consumers may be prepared to accept slightly lower production values in return for access to a much larger variety of music. (They may have no choice if their persistence in using file-sharing systems breaks the existing business models.) Performers and producers may need to find ways to better exploit the possibilities for building on each other’s work. Some have suggested an open-source movement for music, similar to existing open-source software.

17.1.1 Composition Technology

The profits of the modern Western music industry have funded (and encouraged) the development and use of music-related technology. Technology contributes to recording, editing, and distribution; it allows the development of new musical instruments, such as the electric guitar and the synthesizer; and it enables computerised performance, where the musician is replaced by a programmer typing performance instructions into computer software.

But there is one major component of the music production process where technology still plays a very subsidiary role to the efforts of the human musician: composition. There does exist software that can help the musician to compose music, but such software can no more compose good music by itself than a word-processor can write a best-selling novel.

As already discussed in some detail in this book, the non-existence of any effective technology to automate the composition of music is a direct consequence of our lack of understanding of what music is. The process of musical composition remains ad hoc and intuitive. This connection is summed up in Chapter 2 in the “luxury yacht test”: if you knew what music was you would be able to develop an algorithm for musical composition, and from sales of the music composed by the algorithm you would have enough money to buy a luxury yacht.

17.1.2 Profiting from a Complete Theory

The time may come when someone realises how much money could be made from a complete scientific theory of music, and commits a significant investment to the development of such a theory.

A simple business plan is as follows:

- Decide to solve music problem.
- Hire suitable employees and swear them all to secrecy.
- Acquire database of music and perform research thereon.
- Research response to music on human and animal subjects.
- Solve problem sufficiently to algorithmically generate music.
- Use algorithm to compose new strong music.
- Produce and sell the music.
- Rake in profits.

An essential component of this plan is secrecy. Once everyone knows what the composition algorithm is, or enough is known about music to easily derive a composition algorithm, anyone will be able to compose music algorithmically, and the commercial value of individual musical compositions will be reduced to zero.

Given the requirement for secrecy, it is entirely possible that *someone is already carrying out this plan*. It is even possible that multiple parties are carrying out the same plan, each in ignorance of the others.

But unless the release of algorithmically composed music is very carefully controlled, there will be some obvious signs and symptoms of music composed algorithmically:

- Some of the music will be radically distinct from existing genres in ways identifiable even to musically naïve listeners.

- The quality of the music will be noticeably better than music composed the traditional way. This will lead to almost total domination of the commercial music industry by a small number of composers and songwriters.

It is slightly possible that neither of these signs will become apparent:

- It may be that we have already discovered all possible musical genres, so it is not possible to create any new genres.
- It may be that there is some sort of “ceiling” of musical strength, and ad hoc musical composition has already reached this ceiling. Algorithmic composition might increase the number of very strong songs, but the strength of the songs may be no greater than anything in the existing body of music.

17.2 A Post-Music-Theory World

If one person or one group of people can discover the secret of music, then so can others. So even if the first discoverers keep it a secret for commercial reasons, the secret will eventually get out. There are the temptations of fame and credit—once one has made one’s first few millions, one might want to claim the credit for a major scientific discovery before someone else makes the same claim.

The standard legal answer to the problem of commercially exploiting a new discovery which cannot be kept secret is to apply for a patent.¹ The disadvantage of a patent is that it requires disclosure. If the applicant is lucky, there is a gap between initial application (after which commercial exploitation is permitted), and granting of the patent (when disclosure is compulsory). Depending on which country you are in (or more precisely, depending on which country or countries you wish to apply for a patent in), you may or may not be permitted to commercially exploit an invention *before* your application for a patent on the invention. But if a dishonest inventor of a musical composition algorithm secretly used their algorithm to compose music, published the music, received royalties, and then they made a patent application (claiming that they had not yet used the invention commercially), it would be difficult to prove that they had indeed used their invention prior to the application (and therefore were not entitled to receive a patent).

Different countries also vary in whether or not they allow patents on algorithms.

¹Legal Disclaimer: I am not a lawyer. The contents of this book do not constitute legal advice; they are simply my own understanding of the legal situation. If you do solve the music problem, and you are unclear about legal issues, my advice is to get proper legal advice from a lawyer.

Once a patented composition algorithm has been disclosed, preventing infringement may not be easy. For example, it may be possible for a competent software developer to implement the invention with a few hundred lines of code in a high-level programming language. The file-sharing saga has shown that many people will do whatever it takes to get direct and free access to music that they want. The instant pleasure of music can overcome anxieties about illegality, especially if the risk of being caught is low enough.

Whether a musical composition algorithm gets patented or not, the discovery of such an algorithm will have a major impact on the economics of the music industry.

Here is a rough sequence of events that currently happen when someone composes new good quality music:

- The composer composes the music.
- The composer, or a separate lyricist, writes some lyrics, because singing is the preferred form of music for most listeners.
- The song is taken to a publisher.
- The publisher accepts the song, and looks for a performer to perform the song.
- The performer (perhaps already signed up to a record company) signs up to perform the song and make a recording.
- The performer performs and records the song in a recording studio.
- The recorded performance gets mixed by a mixer.
- The mixed recorded performance gets mastered.
- The record company decides to sell the recording.
- Someone makes a video of the performer pretending to sing the same song live against a soundtrack of the mastered recording.
- The record company's promoters promote the song to radio stations and TV music channels.
- The performer undertakes a world tour, playing the new song and any others they happen to have on their new album.
- Consumers hear the song on the radio, see the video and perhaps go to a concert.
- Eventually the song appears in other forms: bands play it in pubs, other well-known performers do cover versions, sheet music becomes available, and lots of people sing it in karaoke bars.

- Consumers have listened to the song so many times that most are moderately bored by it. The song is included in various low cost compilation albums. The original CD appears on sale tables in CD shops.
- Consumers crave their next “fix” of new music.

There are variations on some of these steps, like “the composer also performs the song themselves”, or “few consumers buy the song until a second performer performs a cover version”.

In general the duration from initial composition to the stage of being played excessively on the radio is at least a few months.

Now consider the sequence of events in a post-music-theory world:

- The consumer installs some musical composition software on their computer. It is assumed that they have already installed a good sound-card, good quality speakers and a high quality set of **sound fonts** (i.e. digital versions of instrument sounds).
- The consumer fiddles with input parameters on the composition software, and uses the mouse to click the “Compose” button.
- The software creates a new composition and then plays it through the computer’s sound system.
- If the consumer likes the composition, they save the details, and play it some more.
- If the consumer particularly likes the results of the current composition, they may post it to their web-log, or email it to their friends.
- The consumer gets bored, fiddles a bit more with input parameters for composition, and clicks the “Compose” button again ...

In as much as “consumers” are people who buy things that others have produced, our “consumer” was only really “consuming” when they set up their computer hardware and installed the composition software (and the software might have been free anyway). After that initial step it doesn’t make so much sense to refer to them as a “consumer”, since they are now doing all of their own production.

How long will it take to go from pressing the “Compose” button to hearing the song? If the composition algorithm is computationally intensive, then there might be some delay. In the worst case music lovers might be forced to leave their computers running overnight and see what comes up in the morning. But it is quite possible that gratification will be totally immediate. And the software will be configurable to automate all the above steps: the computer composes a new tune, plays it several times, composes another tune, plays it several times, then repeats a few tunes composed on some

earlier occasion, and so on, with configuration options chosen to match the listener's preferred rate of exposure to new and old music.

There are some things missing from this do-it-yourself (on your own computer) approach to musical composition:

- The quality of performance, mixing and mastering will only be as good as what is programmed into the software, or what the user is able to do for themselves by adjusting options in the software.
- The composition algorithm will probably not generate good lyrics: lyrics are natural language, and generating interesting natural language is a whole separate difficult problem in itself.
- Even if lyrics can be written, the software might not be able to sing very well.
- And even if the software can sing as well as a person, we actually like to know that a person is singing a song to us.

Writing lyrics is a non-trivial skill: if you don't believe this, pick a well-known tune and try writing your own lyrics to it, and see if they sound as good as the original lyrics.

Difficulties with lyrics and singing will create new markets in the music industry, specifically for:

- Lyricists who can write good lyrics quickly.
- Singers who can quickly learn to sing new songs. It may be hard for some singers to do this. (But singers in the new post-music-theory economy will at least be spared the unbearable boredom of having to sing the same hit song over and over and over again at all their concerts.)
- Ditto with instrumentalists, given that there will still be a demand for live performances of music.

17.2.1 Music Junkies?

A cornucopia of music from a composition algorithm based on a complete scientific understanding of music perception may not be an entirely good thing. It may, as the saying goes, be too much of a good thing.

Technology constantly threatens us with new and dangerous addictions: fast cars, television, designer drugs, video games and Internet pornography. Algorithmically generated music may be the next addition to this list. Come home from work, turn on the computer, bring up the software, compose and play some new music. Or just download the latest hot compositions that have been posted on the Internet.

As is the case for many other forms of entertainment, music “addiction” may be self-limiting for most people. If there is some 5 percent of the population (the “music junkies”) whose life is ruined and corrupted by this new pleasure, then so be it, and the rest of us enjoy it in moderation, and life goes on.

We can only hope that the power of rationally composed music over us will not be too great. There might be a campaign to get musical composition algorithms classified as a “Class A” drug, but it seems unlikely that those in power could successfully enforce a law against citizens composing music in the privacy of their own computers.

17.2.2 The Future

Futurology is a difficult enterprise. The things that make the future most interesting are the ideas that become known in the future that were not known in the past. By definition these ideas are not known at the time the prediction is made.

Sometimes even a small technological change has profound consequences for everything, and the full extent of these consequences is not immediately obvious. The intrinsic conservativeness of our thinking makes us reluctant to throw away assumptions about how the world is and how it should be, even if we have observed a change and we know logically that it breaks many of those same assumptions.

Faced with these difficulties, and not wanting to appear too much of a fool to future generations, I will risk just one more prediction about the future of music: the next step in improving our understanding of music may be taken by a reader of this book.