1 of 1 DOCUMENT

New Scientist

August 9, 1997

Requiem for the soul

BYLINE: If creating sublime music is the highest of human achievements, how come a pile of computer code writes better music than most people ? Bob Holmes investigates
Bob Holmes

SECTION: Features, Pg. 2323

LENGTH: 3256 words

MUSIC LOVERS who attended a concert at the University of California at Santa Cruz in April heard the college orchestra put their hearts into Mozart's 42nd symphony. The work features all the glorious harmonies and elegant flourishes one expects from Mozart, though it perhaps lacks some of the genius of other late works such as the 41st symphony, the famous "Jupiter". But there is a reason for this lack of "je ne sais quoi". In the 207-year interval between writing the 41st symphony and the 42nd, Mozart was busy doing more decomposing than composing.

So what exactly is going on ? How could Mozart write a symphony more than 200 years after his death ? Meet a computer program called EMI (pronounced Emmy) and its creator, a living, human composer named David Cope. Under Cope's tutelage, EMI created the 42nd symphony by analysing some of Mozart's other 41 and extracting "essence of Mozart". It used this to generate a brand-new composition in which Mozart nevertheless "wrote" every note and decided every detail of harmony, rhythm and instrumentation. "There's no expert in the world who could, without knowing its source, say for certain that it's not Mozart," says Cope.

EMI, or more properly Experiments in Musical Intelligence, is not just a Mozart specialist, however. In the past few years, it has produced new works by Bach, Beethoven, Brahms, Chopin and the ragtime composer Scott Joplin. It has become a regular collaborator on Cope's own compositions and has even received job offers from the commercial music industry.

The fact that a program running on an ordinary Macintosh computer can produce high-quality music is turning heads in the field of artificial intelligence, as well. "I think it's one of the most provocative, if not the most provocative thing I've come across in artificial intelligence," says Douglas Hofstadter, a cognitive scientist at Indiana University in Bloomington who studies computer creativity.

Hofstadter, a passionate amateur pianist, thinks most of EMI's output still falls short of the real thing. But occasionally it's bang on, as in the case of a "Chopin" mazurka. "When I first played through that mazurka and got to know it, I was quite stunned by it," he says. "It sounded to me, except for a few glitches, as if it could slide right into the book of Chopin mazurkas."

Threatening

Many people - including Hofstadter - find Cope's program profoundly threatening. "EMI has no model whatsoever of life

experiences, has no sense of itself, has no sense of Chopin, has never heard a note of music, has no trace in it of where I think music comes from. Not a trace," he says. "I'm comparing that with an entire human soul, one forged by the struggles and travails of life, and all the experiences that create emotion: turmoil, excitement, hope, despair, resignation, everything you want to think of that goes into building a character."

Yet EMI's mazurka is all but indistinguishable from the real thing. Does that mean, worries Hofstadter, that the composer's soul is irrelevant to the music ? "If that's the case - and I'm not saying it is - then I've been fooled by music all my life. I've been sucked in by a vast illusion. And that would be for me an absolute tragedy, because my entire life I've been moved by music," he says. "I've always felt I've been coming into contact with the absolute essence of humanity."

Cope didn't intend to shake music to its foundations when he began tinkering with the idea of a composing computer. All he wanted was a little help. In 1982, at age 41, Cope was already a successful composer whose work was widely performed. But that year he found himself for the first time facing composer's block. Maybe, he thought, he could devise a program that could look at other music he'd written and suggest which notes might logically come next to fill the void. Fifteen years (and 100 000 lines of computer code) later, that program has become EMI, and Cope is still making improvements.

The basic idea behind EMI came from the "musical dice games" beloved of Mozart and other 18th-century composers. They began with fragments of music that could be played in a random order, and then wrote new works by ordering the fragments according to dice throws ("Dicing with Mozart", "New Scientist", 14 December 1991, p 26). Cope decided to try something similar by taking a familiar piece, breaking it into tiny bits, and then reassembling the bits in a different but still logical order - rather like taking a house apart and using the bricks to build a new house in a similar architectural style.

For most music, however, such slice-and-splice techniques yield gibberish just as surely as shuffling the words in a Hamlet soliloquy. "If you take any Mozart sonata and snip it into measure-length bits and stick them together any old way, it won't sound like Mozart any more," says Cope. "It'll sound like a joke." To make sense of the fragments, EMI needed an understanding of musical structure - a grammar and syntax of music that would keep the elements in logical order.

So Cope devised a way of parsing a work so that he could assign a grammatical function to every musical fragment. Roughly speaking, his method classifies fragments according to whether they begin or end a phrase or serve as decoration. Every phrase, in turn, can be assigned a similar function within larger scale musical form.

As an example of the layered structure of music, consider the tune of "My Bonny Lies Over the Ocean" . The first line (My bonny lies over the ocean) is a beginning, the second (My bonny lies over the sea) is an end, because it "answers" the first. The next two lines (My bonny lies over the ocean/Oh, bring back my bonny to me) form another beginning-end pair. And the first two lines taken together form a beginning at a larger scale, which the last two lines answer to make the end.

EMI, of course, is blind and deaf, so before it can begin to parse a piece of music, Cope must feed in the score in a form it can understand. He translates the music into a strange language of computer-coded "events", which specify for each note when it begins, how long it lasts, its pitch, loudness and which instrument plays it. Then EMI extracts the harmonic structure of the piece - in a generic way that is independent of the actual chords. If "My Bonny lies over the ocean" is played in the key of F, the F chord is called the tonic, while the C chord - which ends the second line - is called the dominant. Alternatively, if the tune was played in G, the tonic would be G and the dominant D. By viewing chords in this way, EMI can find similarities between pieces in different keys.

EMI also analyses the melody - not, however, in terms of notes, but the "intervals" between them (in this case, the number of semitones between notes in a melody). Again, this makes the actual notes played irrelevant but the pattern of intervals all-important.

To get EMI in the mood for composition, Cope feeds it a series of musical examples, say Mozart symphonies. The program then begins to parse the examples. By focusing on cues such as chord changes, melodic patterns and musical context, EMI might divide the first movement of a symphony into four large sections - a main theme, a development in which Mozart varies the theme, a recapitulation, and a closing section, or coda.

The main theme might also divide into two further themes. Within each of these, EMI identifies individual phrases and records their lengths, how they fit together and what key they are in. At the level of individual beats, the program records which instruments

play which notes in a chord and what the next note is for each instrument.

This done, EMI now has a complex catalogue of the choices Mozart made in writing his symphonies. This catalogue takes the form of a set of lists, or lexicons. Each lexicon, says Cope, is "a storehouse of musical bits of the same type". For example, it might hold all the phrase endings in the tonic key that Mozart used in phrases beginning with a "dominant sixth" chord.

To generate a new symphony, EMI picks an opening chord and turns to the appropriate lexicon to ask what Mozart did next. From the range of choices, it selects one at random, then turns to another lexicon to ask what might come next, and so on. All the while, EMI keeps track of larger-scale forms so that musical phrases end when its analysis of the examples says they ought to, and in the correct key. "EMI functions the way most composers function," says Cope. "We struggle with the routine note-to-note rules, but always in the back of our mind we have the idea that we're within this larger piece."

Musical sense

To create the composing side of EMI, Cope borrowed a technique from computer programs that process language: augmented transition networks (ATNs). This provides a way of linking lexicons of words into meaningful sentences. For example, an ATN might link the noun "ball" to the verbs "roll" and "bounce", but not to "slide" or "drink". "Bounce" might link to the prepositions "over" and "into", but not "up", and then to the nouns "wall" and "hand", but not "hair".

This simple ATN - with modifications for articles and possessive pronouns - could yield a variety of different, logical sentences such as "The ball bounced over the wall" or "The ball bounced into my hand" or even "Into my hand the ball bounced". But the ATN excludes such grammatically correct nonsense as "The ball slid up my hair".

EMI uses the same procedure when writing music. As it parses the examples, it constructs ATNs that link together notes, harmonies, rhythms and sets of instruments according to the way Mozart used them. As EMI constructs a new composition, the links within the ATN ensure that what emerges conforms to what made musical sense to Mozart.

By itself, this process creates a workable, but generic-sounding symphony. To capture the distinctive sound that makes Mozart Mozart, Cope found he needed another step. So he grafted on a pattern matcher that sifts through the musical examples looking for short sequences - usually half a dozen or so notes - that show up in piece after piece. These recurring motifs, which Cope calls "signatures", represent the little, mostly unconscious turns of phrase that make each composer's work distinctive to the ear, just as we might recognise a friend's speech by his habit of saying "basically" a lot (see Diagrams).

Instead of breaking these signatures into their component notes, EMI compiles them into a separate list. It then drops them into a composition at points where the ATN indicates that Mozart might have used them, so any listener with a little musical background will say, "Yes, that's Mozart," or "That sounds like Chopin". Curiously, this rough-and-ready pattern-matcher seems to capture an important part of a human composer's essence.

So armed, EMI can turn out credible music in styles that are always recognisable, even if the music itself may be a cut below the level of the masters. "It definitely sounds like Rachmaninoff," says Anatole Leikin, a pianist at the University of California at Santa Cruz who has performed a "Rachmaninoff" piano work created by EMI. "But I would say it's someone who imitates Rachmaninoff - a composer of lesser stature."

Even Cope concedes that EMI's music usually lacks the spark of true genius - a shortcoming that leaves many listeners dissatisfied. But then again, the music of most human composers also lacks genius. "As far as I'm concerned, EMI 's Mozart is better than 99 per cent of non-Mozart classical music," he says. Cope compares EMI favourably with the best of Antonio Salieri, Mozart's contemporary and - if Peter Shaffer's play "Amadeus"is to be believed - nemesis. "His music is not very good," says Cope. "This music is better than that."

What about the musicians who perform EMI's works ? Can they feel the same depth, richness and emotion in the music issuing from EMI's carefully constructed intelligence as they can in music from a living, breathing, hurting human ? Most performers say yes, at least some of the time. Of course, as Leikin points out, since EMI is a computer program it doesn't actually have any emotions to put into its music. "The question is whether we can "hear" it emotionally," he says. "My experience is we can, at least in certain

sections."

Certainly, EMI has proven skilled enough to fill the role it was designed for, that of composer's helpmate. It offers fresh alternatives whenever Cope hits a difficult spot in a new work. And, having studied its master's music, compiled lexicons from his work and found his signatures, the program suggests options that fit with his style. The composer might use EMI's suggestion verbatim, adapt it, or ignore it entirely. But having the choice can be a great comfort. "It's how I think composers in the 21st century will work," says Cope. "All the Cope pieces I write now are done with EMI. I will never write another piece as just David Cope."

When he is not using EMI to help with his own compositions, Cope sets it loose on his wish list of new works by dead composers. "I love Mahler, but I'm so tired of the 10 symphonies. God, it would be great to hear a new one," he says. Cope reckons Mahler's 11th will take two or three years to complete. It's number two on the list tacked to the wall of his studio. Above it is the Mozart symphony performed in April - EMI's first full symphony to be performed in public. Below it are eight other symphonies by Rachmaninoff, Brahms, Beethoven and others. "The Mozart is done, so there's one down and nine to go."

Like all good mentors, Cope has also given EMI the chance to develop its own style. In 1992, he started it off with three pieces by the highly original 20th-century composer Igor Stravinsky. As EMI generated new pieces in the Russian's style, Cope began to replace the Stravinsky examples with EMI's own output. He continued switching, using EMI's most recent compositions as examples for its next ones, so that EMI's style drifted further and further from Stravinsky's. Now and then, he also threw into the mix a piece of music from another composer of the same era. "I was pretending EMI was a human," says Cope. "Humans don't just compose in a vacuum, and I didn't want EMI to, either. I tried to make EMI more cosmopolitan."

In three months of continuous computing, EMI churned out 5000 of its own works, including 1500 symphonies, 2000 piano sonatas and 1500 miscellaneous pieces - far more than Cope could possibly listen to in his lifetime. But as he browses through this massive output, he is finding some interesting music. He is now, for example, transcribing EMI's 1383rd symphony, in which he hears echoes of a Mahler work that EMI had studied a short time earlier. "I'm marvelling," says Cope. "It's quite good at times. Other parts confuse the hell out of me as to why it's doing what it's doing." One thing is clear, however: in its three months of brooding introspection, EMI developed its own individual style, which Cope describes as that of a unique Russian-American composer working between 1918 and 1935.

EMI's ability to go beyond mimicking actual composers opens the door to some fascinating what-if games. What if Mozart had lived a decade longer and known Beethoven, for example ? Until EMI, no one could do more than idly speculate. But now, we could simply add some well-chosen Beethoven to the program's Mozart database and listen to one version of what might have happened. Or if, in a fit of musical eugenics, one threw together the best of Bach, Beethoven and Brahms, would this breed a new classical supercomposer or a grotesque freak ?

"All those scary questions are out there. They've never been asked in music before, because music's so staid and circumscribed," says Eleanor Selfridge-Field, a musicologist at Stanford University, California. Indeed, few AI specialists have ever ventured so far into the realm of art and aesthetics. Selfridge-Field is familiar with just a handful of other programs that try to simulate musical style. "Their results have been only a small fraction of what David's accomplished," she says.

Hofstadter, too, thinks EMI stands almost alone in the field of artistic artificial intelligence. But though the subject matter is unusual, he says, Cope's approach sits squarely in the mainstream with its reliance on sheer processing power rather than elegant reasoning or machine learning. In that respect, EMI is a close cousin to Deep Blue, IBM's chess-playing program which recently defeated Garry Kasparov. "EMI does a good job, but it involves a very brute force, non-plausible model of what humans do, just as I think Deep Blue is a very implausible model of cognitive processes," says Hofstadter.

Successful as it is, EMI is not yet good enough to convincingly mimic such musical thoroughbreds as Mozart and Beethoven every time. But the carthorses of the musical world - the writers of advertising jingles and B-movie scores - should prove much easier to mimic. If so, these composers may soon find machines muscling in on their jobs, just as synthesisers began replacing human instrumentalists in the 1980s. Indeed, Cope says he has already been approached by people who are interested in hiring EMI to produce music for commercial uses.

Yet despite the potential riches that EMI could bring, Cope has not sought to patent it. Quite the contrary: much of the computer

code for EMI is published, and he will send the entire program to anyone who wants it. Cope is reluctant, however, to share his experience of getting the best out of EMI. "It would horrify me personally to walk into a hall and hear someone else's EMI Mozart symphony and hear that it was as good as one I produced," he says.

Taking the credit

Cope's experience turns out to be a crucial part of EMI's success, because it is Cope who decides which examples EMI should analyse before assembling a new piece - a choice that greatly affects the mood of the finished work. And when gathering signatures for a composition, Cope often tinkers with the pattern matcher for hours, tweaking one and then another of its 22 different controls until the signatures sound as though they were really written by Bach or Mozart.

"This is a very personal operation," says Cope. "I definitely want to take credit for most of the good work that EMI has produced in the sense that I very much had a hand in developing the signatures."

And, most important of all, Cope decides which pieces are good enough to let out of his studio. Most of EMI's music goes straight in the bin once Cope has picked over it. The small minority of pieces that survive are those that Cope finds pleasing - a judgment that EMI cannot make for itself.

Ironically, the computer program that sometimes produces music as sublime as Mozart's can't tell the difference between a work of genius and a piece of lift music. In the end, that aesthetic judgment may prove to be the one thing that no programmer can code for.

Further Reading: More information about Cope's work can be found in his books: "Computers and Music Style" and "Experiments in Musical Intelligence", both published by A-R Editions. To hear some of EMI's work, go to the listening booth at http://arts.ucsc.edu/faculty/cope/home.

For more science news see http://www.newscientist.comMagazine

LOAD-DATE: August 8, 1997

LANGUAGE: EnglishFeatures

Copyright 1997 Reed Business Information UK, a division of Reed Elsevier, Inc.
All Rights Reserved.