Music from Inverted Analysis: Creativity or Just Copying?

Chris Thornton

October 30, 2009

Abstract

Although it generates excellent compositions in the style of Mozart, Chopin and other composers, David Cope's EMI system is often viewed as a fancy copying machine rather than a system which shows genuine creativity. And it is true that the generative functionality of the system (in its best-known form) relies on careful, hand-modulated modelling of examples drawn from a large database of human compositions. But how well would an EMI-like system do with less human intervention? In the talk I'll introduce and demo a system I've developed which examines the issue. The program uses analytic inversion (like EMI) but eliminates all other aspects of manual intervention, relying entirely on generic modelling and instantiation methods. The results achieved so far are not that impressive so I'm keen to hear suggestions about how to take the work forward. If you haven't heard any EMI output before, you can play examples at Cope's MP3 website (`http://arts.ucsc.edu/faculty/cope/experiments.htm`) I'll also be playing a couple of EMI compositions in the talk.

Artefact Generation from Inverted Analysis

file called creativity-from-analysis-inversion

Early Cope

In his work with the EMI system (Experiments in Musical Intelligence), David Cope has explored many uses of inverted analysis for music generation.

Basic approach uses 'recombinance' with 'voice leading' and 'texture matching'.

Assemble a database of compositions, chop them up into musically meaningful units and recombine so as to enforce voice and texture continuities.

Results may be quite good 'locally' but often lack higher-level structure.

Later Cope

Cope has also explored ways of enforcing higher-level structures (mediated by allusions, quotations, paraphrases, likenesses, frameworks, commonalities) at both analytic and generative stages. Also how to detect them using learning/analogy methods etc.

Also (and particularly) ways of modulating these processes using SPEAC coding.

Represent music as a standard sequence of abstract, musical events: S=statement, P=preparation, E=extension, A=antecedant, C=consequent. Allow hierarhies and fuzzy categorizations.

Structures rules (from p. 237-8 of 2005).

```
S --> P S | A C | S E
P --> P E | E E | A C
E --> S E | E E
```

Output

http://arts.ucsc.edu/faculty/cope/mp3page.htm

Latest Cope

Cope has tended to apply methods to fine-tuned databases containing works of noted composers (Mozart, Stravinsky, Chopin, Bach, Beethoven...)

Latest Cope

Cope has tended to apply methods to fine-tuned databases containing works of noted composers (Mozart, Stravinsky, Chopin, Bach, Beethoven...)

- Good news: output can sound just like new versions of the 'real thing'

Latest Cope

Cope has tended to apply methods to fine-tuned databases containing works of noted composers (Mozart, Stravinsky, Chopin, Bach, Beethoven...)

- Good news: output can sound just like new versions of the 'real thing'
- Bad news: people get upset! Cope's compositions can sound highly derivative.

Latest Cope

Cope has tended to apply methods to fine-tuned databases containing works of noted composers (Mozart, Stravinsky, Chopin, Bach, Beethoven...)

- Good news: output can sound just like new versions of the 'real thing'
- Bad news: people get upset! Cope's compositions can sound highly derivative.

Because Cope wants to produce good music (and feels that his system just does what he would do but quicker) he views this reaction as a problem.

Now aims to persue similar methods but without using 'historical databases.'

New system to be called 'Emily Howell'

But has he over-reacted? The problem was not the use of

How about a Cope-free emmy?

Although EMI relies on (forms of) analytic-inversion, Cope has added a vast assembly of interacting knowledge, mechanism and hand-modulation.

It is hard to make any analysis of the value of any particular component.

So as an experiment...

Build a minimal version of the system which eliminates everything except for the inner loop.

No labeling, no SPEAC, no voice leading, no texture matching, no paraphrases, no quotations, no signatures, no association net, no learning, no analogy detection...

Just use *generic* sequence analysis and instantiation.

A simple way of analysing the structure of sequential data.

Work out frequencies with which possible 'n-grams' (i.e., sub-sequences of n items) occur in the data.

A new sequence with the same (sequential) statistical properties can then be generated by 'Markov chaining'.

Start with some random seed element and keep adding probable continuations.

Markov chaining

Sequence: A B D A D D A A B

Calculate frequencies for all sub-sequences of length 2.

```
'A B' 2
'B D' 1
'D A' 2
'A D' 1
'D D' 1
'A A' 1
```

Now generate a new sequence starting with B.

Markov chaining

Sequence: A B D A D D A A B

Calculate frequencies for all sub-sequences of length 2.

```
'A B' 2
'B D' 1
'D A' 2
'A D' 1
'D D' 1
'A A' 1
```

Now generate a new sequence starting with B.

- Add D (p=1.0)

Markov chaining

Sequence: A B D A D D A A B

Calculate frequencies for all sub-sequences of length 2.

```
'A B' 2
'B D' 1
'D A' 2
'A D' 1
'D D' 1
'A A' 1
```

Now generate a new sequence starting with B.

- Add D (p=1.0)
- Add A (p=0.66) or D (p=0.33)

Markov chaining

Sequence: A B D A D D A A B

Calculate frequencies for all sub-sequences of length 2.

```
'A B' 2
'B D' 1
'D A' 2
'A D' 1
'D D' 1
'A A' 1
```

Now generate a new sequence starting with B.

- Add D (p=1.0)
- Add A (p=0.66) or D (p=0.33)
- If added A, Add B (p=0.65), A (p=0.25), or D (p=0.25)

Sequence: A B D A D D A A B

Calculate frequencies for all sub-sequences of length 2.

```
'A B' 2
'B D' 1
'D A' 2
'A D' 1
'D D' 1
'A A' 1
```

Now generate a new sequence starting with B.

- Add D (p=1.0)
- Add A (p=0.66) or D (p=0.33)
- If added A, Add B (p=0.65), A (p=0.25), or D (p=0.25)
- If added D, Add A (p=0.66, D (p=0.33)

Markov chaining

Sequence: A B D A D D A A B

Calculate frequencies for all sub-sequences of length 2.

```
'A B' 2
'B D' 1
'D A' 2
'A D' 1
'D D' 1
'A A' 1
```

Now generate a new sequence starting with B.

- Add D (p=1.0)
- Add A (p=0.66) or D (p=0.33)
- If added A, Add B (p=0.65), A (p=0.25), or D (p=0.25)
- If added D, Add A (p=0.66), D (p=0.33)

Markov chaining

Sequence: A B D A D D A A B

Calculate frequencies for all sub-sequences of length 2.

```
'A B' 2
'B D' 1
'D A' 2
'A D' 1
'D D' 1
'A A' 1
```

Now generate a new sequence starting with B.

- Add D (p=1.0)
- Add A (p=0.66) or D (p=0.33)
- If added A, Add B (p=0.65), A (p=0.25), or D (p=0.25)
- If added D, Add A (p=0.66), D (p=0.33)

Markov Beatles

Markov chains generated from Beatles lyrics using different size analyses.

Markov chains generated from Beatles lyrics using different size analyses.

- n = 1

Markov Beatles

Markov chains generated from Beatles lyrics using different size analyses.

- n = 1

you the to you don't my i a of could and i that to you the should me come you i you i the been to you the she and to you i the you it but you it you i she's the i the i i you please my

Markov Beatles

Markov chains generated from Beatles lyrics using different size analyses.

- n = 1

you the to you don't my i a of could and i that to you the should me come you i you i the been to you the she and to you i the you it but you it you i she's the i the i i you please my

- n = 2

Markov Beatles

Markov chains generated from Beatles lyrics using different size analyses.

- n = 1

you the to you don't my i a of could and i that to you the should me come you i you i the been to you the she and to you i the you it but you it you i she's the i the i i you please my

- n = 2

you can see the girl when you can show but it's getting so many years yes, wait till tomorrow way get back in my apart but i know we will love you know my life i've never be mine, i want that's why i know i know i held each

Markov Beatles

Markov chains generated from Beatles lyrics using different size analyses.

- n = 1

you the to you don't my i a of could and i that to you the should me come you i you i the been to you the she and to you i the you it but you it you i she's the i the i i you please my

- n = 2

you can see the girl when you can show but it's getting so many years yes, wait till tomorrow way get back in my apart but i know we will love you know my life i've never be mine, i want that's why i know i know i held each

- n = 3

Markov chains generated from Beatles lyrics using different size analyses.

- n = 1

you the to you don't my i a of could and i that to you the should me come you i you i the been to you the she and to you i the you it but you it you i she's the i the i i you please my

- n = 2

you can see the girl when you can show but it's getting so many years yes, wait till tomorrow way get back in my apart but i know we will love you know my life i've never be mine, i want that's why i know i know i held each

- n = 3

you i want me to dance with you and i feel as though you ought to do what he left it won't be my baby, now went wrong i've got a boat on me and so my name you know my baby everybody's trying to be a boat on the hill sees the sheik of love chains of love, come

To discover hierarchical sequence structure, n-gram analysis must be performed hierarchically and permit disjunctive constructions.

To discover hierarchical sequence structure, n-gram analysis must be performed hierarchically and permit disjunctive constructions.

(1) Perform ordinary (disjunct-allowing) n-gram analysis

Sequence reduction

To discover hierarchical sequence structure, n-gram analysis must be performed hierarchically and permit disjunctive constructions.

(1) Perform ordinary (disjunct-allowing) n-gram analysis
(2) Replace sequence with its shortest labelling in terms of identified n-grams

Sequence reduction

To discover hierarchical sequence structure, n-gram analysis must be performed hierarchically and permit disjunctive constructions.

(1) Perform ordinary (disjunct-allowing) n-gram analysis
(2) Replace sequence with its shortest labelling in terms of identified n-grams
(3) If the sequence has more than one item, go back to step 1

Sequence reduction

To discover hierarchical sequence structure, n-gram analysis must be performed hierarchically and permit disjunctive constructions.

(1) Perform ordinary (disjunct-allowing) n-gram analysis
(2) Replace sequence with its shortest labelling in terms of identified n-grams
(3) If the sequence has more than one item, go back to step 1

Produces a *hierarchical* representation of sequence structure.

Example

Analysis of ‘the cat sat on the cat’

A	‘the’	2
B	‘cat’	2
C	‘sat’	1
D	‘on’	1
E	‘the cat’	2
F	‘cat sat’	1
G	‘sat on’	1
H	‘on the’	1

Hierarchical analysis

```
the cat sat on the cat
                (n-grams as above)
```

Generation

Hierarchical structure derived this way can be used (like a phrase-structure grammar) to generate new sequences with the same statistical properties.

Given

```
K --> I E
I --> E G
E --> the cat
G --> sat on
```

We recursively expand top label.

```
K ==> I E
I E ==> E G the cat
E G the cat ==> the cat sat on the cat
```

MuSeR = Music from Sequence Reduction

Generic/hierarchical/disjunctive sequence analysis/generation with musical sequences.
archive=SequenceReduction.jar height

Evaluation

Given the terrible performance, the indication is that analytic inversion plays an insignificant role in EMI.

But it may be possible to get the performance closer to the EMI level.

Evaluation

Given the terrible performance, the indication is that analytic inversion plays an insignificant role in EMI.

But it may be possible to get the performance closer to the EMI level.

- Use bigger (or more finely tuned) ‘databases’

Evaluation

Given the terrible performance, the indication is that analytic inversion plays an insignificant role in EMI.

But it may be possible to get the performance closer to the EMI level.

- Use bigger (or more finely tuned) 'databases'
- Change sequence encoding?

Evaluation

Given the terrible performance, the indication is that analytic inversion plays an insignificant role in EMI.

But it may be possible to get the performance closer to the EMI level.

- Use bigger (or more finely tuned) ‘databases’
- Change sequence encoding?

Present approach uses pitch-offset encoding (which implicitly abstracts key) but absolute duration encoding.

This may mean tempo differences between pieces are obscuring useful generalizations.

Switch to duration-change encoding? Would this work?

Any other ideas????

References

To play with the MUSER applet, follow ‘demos’ from www.christhornton.eu.

Comments welcome (particularly on what causes the system to lock up)

Cope’s web site has (bits and pieces of) his own software.
`http://arts.ucsc.edu/faculty/cope/`

Hofstadter, D. (2002). Staring emmy straight in the eye — and doing my best not to flinch. In T. Dartnall (Ed.), CREATIVITY, COGNITION AND KNOWLEDGE: AN INTERACTION (pp. 67-104). London: Praeger.

Cope, D. (2001). VIRTUAL MUSIC: COMPUTER SYNTHESIS OF MUSICAL STYLE. London: The MIT Press.

Cope, D. (2005). COMPUTER MODELS OF MUSICAL CREATIVITY. The MIT Press.