

Music-making with Hierarchical Markov Models

Chris Thornton
COGS/Informatics
University of Sussex
Brighton
BN1 9QH
UK
c.thornton@sussex.ac.uk

November 13, 2009

In the emerging field of empirical creativity, new artifacts (e.g., drum rhythms) are generated from models culled from one or more existing artifacts. Applied to music, this approach often involves use of Markov models. But these may be insensitive to the global, hierarchical properties that are particularly significant for music. The talk introduces the *hierarchical* Markov model and shows how it enables both large-scale and small-scale properties to be captured in a uniform way. It also presents some examples illustrating application of the approach to melodic sequences.

Empirical creativity (EC)

Empirical creativity can be defined as 'the generation of variations of an existing artefact, using a model of that artefact.'

Basic 1-to-1 approach:

Empirical creativity (EC)

Empirical creativity can be defined as 'the generation of variations of an existing artefact, using a model of that artefact.'

Basic 1-to-1 approach:

- Choose an artefact

Empirical creativity (EC)

Empirical creativity can be defined as 'the generation of variations of an existing artefact, using a model of that artefact.'

Basic 1-to-1 approach:

- Choose an artefact
- Build a model for it.

Empirical creativity (EC)

Empirical creativity can be defined as 'the generation of variations of an existing artefact, using a model of that artefact.'

Basic 1-to-1 approach:

- Choose an artefact
- Build a model for it.
- Use the model to generate one or more variations.

Empirical creativity (EC)

Empirical creativity can be defined as 'the generation of variations of an existing artefact, using a model of that artefact.'

Basic 1-to-1 approach:

- Choose an artefact
- Build a model for it.
- Use the model to generate one or more variations.

Many-to-1 approach:

Empirical creativity (EC)

Empirical creativity can be defined as 'the generation of variations of an existing artefact, using a model of that artefact.'

Basic 1-to-1 approach:

- Choose an artefact
- Build a model for it.
- Use the model to generate one or more variations.

Many-to-1 approach:

- Choose several artifacts.

Empirical creativity (EC)

Empirical creativity can be defined as 'the generation of variations of an existing artefact, using a model of that artefact.'

Basic 1-to-1 approach:

- Choose an artefact
- Build a model for it.
- Use the model to generate one or more variations.

Many-to-1 approach:

- Choose several artifacts.
- Build models for all of them.

Empirical creativity (EC)

Empirical creativity can be defined as 'the generation of variations of an existing artefact, using a model of that artefact.'

Basic 1-to-1 approach:

- Choose an artefact
- Build a model for it.
- Use the model to generate one or more variations.

Many-to-1 approach:

- Choose several artifacts.
- Build models for all of them.
- Combine them to produce a new artefact

Empirical creativity (EC)

Empirical creativity can be defined as 'the generation of variations of an existing artefact, using a model of that artefact.'

Basic 1-to-1 approach:

- Choose an artefact
- Build a model for it.
- Use the model to generate one or more variations.

Many-to-1 approach:

- Choose several artifacts.
- Build models for all of them.
- Combine them to produce a new artefact

Other: modality mash-ups (text-music hybrids), meta composition tools, super-instruments...

Examples

Statistical modeling has often been used to generate text approximations.

Examples

Statistical modeling has often been used to generate text approximations.

- From (Shannon, 1948, p. 7), third-order, character approximation of English text:

Statistical modeling has often been used to generate text approximations.

- From (Shannon, 1948, p. 7), third-order, character approximation of English text:

IN NO IST LAT WHEY CRATICT FROURE BIRS GROCID
PONDENOME OF DEMONSTURES OF THE REPTAGIN IS
REGOACTIONA OF CRE.

Examples

Statistical modeling has often been used to generate text approximations.

- From (Shannon, 1948, p. 7), third-order, character approximation of English text:

IN NO IST LAT WHEY CRATICT FROURE BIRS GROCID
PONDENOME OF DEMONSTURES OF THE REPTAGIN IS
REGOACTIONA OF CRE.

- Second-order word approximation:

Statistical modeling has often been used to generate text approximations.

- From (Shannon, 1948, p. 7), third-order, character approximation of English text:

IN NO IST LAT WHEY CRATICT FROURE BIRS GROCID
PONDENOME OF DEMONSTURES OF THE REPTAGIN IS
REGOACTIONA OF CRE.

- Second-order word approximation:

THE HEAD AND IN FRONTAL ATTACK ON AN ENGLISH
WRITER THAT THE CHARACTER OF THIS POINT IS
THEREFORE ANOTHER METHOD FOR THE LETTERS THAT
THE TIME OF WHO EVER TOLD THE PROBLEM FOR AN
UNEXPECTED

Beatles lyrics with Markov models

Beatles lyrics with Markov models

- 1st-order word approx

Beatles lyrics with Markov models

- 1st-order word approx

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

Beatles lyrics with Markov models

- 1st-order word approx

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

- 3rd-order word approx

Beatles lyrics with Markov models

- 1st-order word approx

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

- 3rd-order word approx

you want to dance with me i'm in love with another oh, when i kiss
her majesty's a pretty nice girl but i'm miles above you tell me, i'm
so alone don't bother me i'm a loser i'm a loser and i'm not a
second time that was so long bye, bye, bye. lady madonna children
at your man i wanna be your man i wanna be let it be, you know
she thinks of him ...

Markov music approximation

demo Markov approximations of Bach prel #1

The problem of large-scale structure

Using typical modeling methods (e.g., HMMs), empirical approximations of (note-represented) music often 'wander'.

We get good reproduction of local patterns but poor or non-existent reproduction of larger-scale or hierarchical structure.

Need to capture both small-scale and large-scale structure in a way that's generatively effective.

This should not involve use of domain theory, on pain of obstructing genre-crossing

Currently exploring hierarchical Markov models 'HiMMS'

Refinement definition

Though normally used for analysis of signaling, information theory can also be used to model informational refinement of symbolic data.

Given a symbolic dataset, we can substitute any construct with a choice that includes that construct.

E.g., from the string 'a b c' we might produce 'a b/d c'.

Each such 'uncertainization' produces

Refinement definition

Though normally used for analysis of signaling, information theory can also be used to model informational refinement of symbolic data.

Given a symbolic dataset, we can substitute any construct with a choice that includes that construct.

E.g., from the string 'a b c' we might produce 'a b/d c'.

Each such 'uncertainization' produces

- (1) a loss of aggregate information, but potentially also

Refinement definition

Though normally used for analysis of signaling, information theory can also be used to model informational refinement of symbolic data.

Given a symbolic dataset, we can substitute any construct with a choice that includes that construct.

E.g., from the string 'a b c' we might produce 'a b/d c'.

Each such 'uncertainization' produces

- (1) a loss of aggregate information, but potentially also
- (2) a reduction in the number of symbols.

Refinement definition

Though normally used for analysis of signaling, information theory can also be used to model informational refinement of symbolic data.

Given a symbolic dataset, we can substitute any construct with a choice that includes that construct.

E.g., from the string 'a b c' we might produce 'a b/d c'.

Each such 'uncertainization' produces


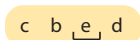
- (1) a loss of aggregate information, but potentially also
- (2) a reduction in the number of symbols.

If mean symbol information (MSI) rises, the data are 'informationally enriched'. Otherwise, they are 'depleted'.

Informational refinement = enriching uncertainization.

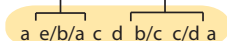

Sequence uncertainization illustrations

Base symbol information = 4 bits





-1

$$\frac{\text{inf} \quad \text{loss} \quad \text{MSI}}{(16.0 - 1.0) / 4 = 3.75}$$



-1.58 -1 -1

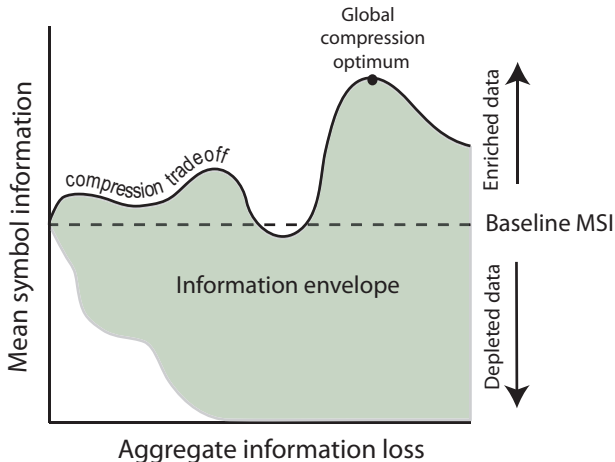
$$\frac{\text{inf} \quad \text{loss} \quad \text{MSI}}{(28.0 - 4.58) / 7 = 3.35}$$



-1 -1

$$\frac{\text{inf} \quad \text{loss} \quad \text{MSI}}{(36.0 - 2.0) / 7 = 4.86}$$

The information envelope



At max concentration, tradeoff between size and accuracy is optimized and we have optimal exploitation of 'patterns'.

Formalization (optional)

Consider a sequence s to be a symbolic specification for a sequence of distributions.

So 'x y z' seen as specifying distributions

$$P(x)=1.0 \quad P(y)=0.0 \quad P(z)=0.0$$

$$P(x)=0.0 \quad P(y)=1.0 \quad P(z)=0.0$$

$$P(x)=0.0 \quad P(y)=0.0 \quad P(z)=1.0$$

Total information content of s :

$$I(s) = \sum_i \log |s_i| - H(s_i)$$

where s_i denotes i 'th distribution.

Mean information

Mean symbol information (MSI):

$$\bar{I}(s) = \frac{I(s)}{|s|}$$

Sequence u^s *uncertainizes* s if it specifies the same sequence of distributions with greater total entropy.

Information loss produced by an uncertainization:

$$L(u^s) = \sum_i H(U_i^s)$$

Mean symbol information in an uncertainization:

$$\bar{I}(U^s) = \frac{I(s) - L(U^s)}{|U^s|}$$

Refinement models

Refinement can be done recursively.

While (some dataset) D contains more than one symbol:

Refinement models

Refinement can be done recursively.

While (some dataset) D contains more than one symbol:

- generate refinement of D .

Refinement can be done recursively.

While (some dataset) D contains more than one symbol:

- generate refinement of D .
- set the refinement to be the new D and repeat.

Refinement can be done recursively.

While (some dataset) D contains more than one symbol:


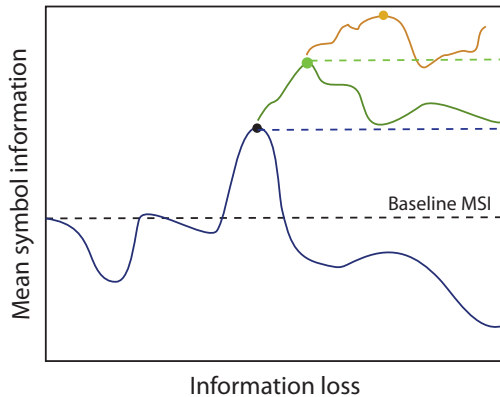
- generate refinement of D .
- set the refinement to be the new D and repeat.

This produces '1st refinement', '2nd refinement', '3rd refinement' etc.

Each choice generalizes over symbols in prior refinement.

The overall tree structure ('refinement model') models small-scale, large-scale and hierarchical properties in the data.

Information-structure derivation



Recursive refinement formalization (optional)

Define the informationally optimal uncertainizations of s as

$$\dot{U}^s = \operatorname{argmax}_{u \in U(S)} \bar{I}(u)$$

First refinement

$$r(s) = \operatorname{argmax}_{u \in \dot{U}(S)} \bar{I}(u)$$

Second refinement:

$$r(r(s))$$

Third refinement:

$$r(r(r(s)))$$

'eg1' demo

Refinement models can be used generatively through a kind of top-down symbol expansion.

At each choice point ('disjunctive branch'), selection can be made at random or in a way that is mediated by the distribution(s) associated with the choice.

Can also learn a separate Markov model for each refinement and use this to shape choices, with distributions being calculated against existing portion of reproduction.

Refinements are hierarchically related, so Markov models are too.

Refinement model becomes a HiMM.

Top-down symbol expansion becomes a 'Markov cascade'.

Replexes and hyper-replexes

Refinement models of existing music can be used to generate informational variations of those pieces.

But don't want to look like a Cope-esque plagiarist...

Make it a formally-defined process.

Replexes and hyper-replexes

Refinement models of existing music can be used to generate informational variations of those pieces.

But don't want to look like a Cope-esque plagiarist...

Make it a formally-defined process.

- X defined to be a 'replex' of Y, if X is generated from a refinement model for Y.

Replexes and hyper-replexes

Refinement models of existing music can be used to generate informational variations of those pieces.

But don't want to look like a Cope-esque plagiarist...

Make it a formally-defined process.

- X defined to be a 'replex' of Y, if X is generated from a refinement model for Y.
- X defined to be a 'hyper-replex' of Y, Z etc. if X is generated from combinations of refinement models for Y, Z etc.

Replexes and hyper-replexes

Refinement models of existing music can be used to generate informational variations of those pieces.

But don't want to look like a Cope-esque plagiarist...

Make it a formally-defined process.

- X defined to be a 'replex' of Y, if X is generated from a refinement model for Y.
- X defined to be a 'hyper-replex' of Y, Z etc. if X is generated from combinations of refinement models for Y, Z etc.

Approximations generated from models that are *partially* enriching are then *partial* replexes.

Can define an optimal refinement model as one that has optimal MSI throughout.

Must offer a maximally compact, 'endogenous' representation at each level.

A reproduction culled from an optimal refinement model is a 'perfect replex'.

At each level, any excess of MSI over baseline constitutes evidence of representational redundancy.

A fully enriching refinement model then represents *redundancy structure*.

Text replexes 'simpsons1' and 'george1'

Music replexes and hyper-replexes

Summary and questions

Summary and questions

- Empirical creativity with symbolic music needs models that capture both large-scale and small-scale properties.

Summary and questions

- Empirical creativity with symbolic music needs models that capture both large-scale and small-scale properties.
- Refinement models augmented with level-specific Markov models (i.e., HiMMs) seem to have the right properties and work fairly well in practice.

Summary and questions

- Empirical creativity with symbolic music needs models that capture both large-scale and small-scale properties.
- Refinement models augmented with level-specific Markov models (i.e., HiMMs) seem to have the right properties and work fairly well in practice.
- HiMMs derived from text often look like probabilistic grammars. Any linguistic implications?

Summary and questions

- Empirical creativity with symbolic music needs models that capture both large-scale and small-scale properties.
- Refinement models augmented with level-specific Markov models (i.e., HiMMs) seem to have the right properties and work fairly well in practice.
- HiMMs derived from text often look like probabilistic grammars. Any linguistic implications?
- Might HiMMs have applications in other sequence modeling tasks?

Summary and questions

- Empirical creativity with symbolic music needs models that capture both large-scale and small-scale properties.
- Refinement models augmented with level-specific Markov models (i.e., HiMMs) seem to have the right properties and work fairly well in practice.
- HiMMs derived from text often look like probabilistic grammars. Any linguistic implications?
- Might HiMMs have applications in other sequence modeling tasks?

Thornton, C. (2009). Representation recovers information. *Cognitive Science*, 33, No. 8 (pp. 1-30).

Thornton, C. (2009). Hierarchical markov modeling for generative music. In G. Scavone, V. Verfaillie and A.D. Silva (Eds.), *Proceedings of the International Computer Music Conference ICMC-09* (pp. 49-52). Montreal, Canada: The International Computer Music Association.

Thornton, C. (2009). Self-redundancy in music. In J. McCormack, M. D'Inverno and M. Boden (Eds.), *Proceedings of the 2009 Dagstuhl Seminar on Computational Creativity: An Interdisciplinary Approach*.

The 'replexer' applet builds refinement models using a combination of heuristic, hill-climbing search and dynamic programming.

Can be used to generate music, text replexes and and music/text

- Shannon, C. (1948). A mathematical theory of communication.
Bell System Technical Journal, 27 (pp. 379-423 and 623-656).
- Thornton, C. (2009). Hierarchical markov modeling for generative music. In G. Scavone, V. Verfaillie and A.D. Silva (Eds.),
Proceedings of the International Computer Music Conference ICMC-09 (pp. 49-52). Montreal, Canada: The International Computer Music Association.