

Statistical Modelling of Music

Christophe Rhodes

Goldsmiths, University of London

Tuesday 3rd April

Introduction

Motivation

Modelling

Chords
Structure

Future

Summary

Music: a set of complicated interactions between

- composer
- arranger
- producer
- performer
- listener
- years of enculturation

Also: an enjoyable activity!

Music: a set of complicated interactions between

- composer
- arranger
- producer
- performer
- listener
- years of enculturation

Also: an enjoyable activity!

Introduction

Motivation

Modelling

Chords
Structure

Future

Summary

Music: a set of complicated interactions between

- composer
- arranger
- producer
- performer
- listener
- years of enculturation

Also: an enjoyable activity!

Introduction

Motivation

Modelling

Chords
Structure

Future

Summary

Music: a set of complicated interactions between

- composer
- arranger
- producer
- performer
- listener
- years of enculturation

Also: an enjoyable activity!

Two reasons:

- to build tools to achieve specific tasks;
- to increase our understanding of musical processes.

Don't (yet) model the entire world: work on smaller aspects.

Two reasons:

- to build tools to achieve specific tasks;
- to increase our understanding of musical processes.

Don't (yet) model the entire world: work on smaller aspects.

Two reasons:

- to build tools to achieve specific tasks;
- to increase our understanding of musical processes.

Don't (yet) model the entire world: work on smaller aspects.

Two reasons:

- to build tools to achieve specific tasks;
- to increase our understanding of musical processes.

Don't (yet) model the entire world: work on smaller aspects.

The task: identifying chords and assigning labels in music

- currently to MIDI transcriptions of performances;
- could be applied to audio directly (given suitable processing).

Applications:

- generating fake books, guitar chords
- feeding into models of music cognition, melodic memory

The task: identifying chords and assigning labels in music

- currently to MIDI transcriptions of performances;
- could be applied to audio directly (given suitable processing).

Applications:

- generating fake books, guitar chords
- feeding into models of music cognition, melodic memory

The task: identifying chords and assigning labels in music

- currently to MIDI transcriptions of performances;
- could be applied to audio directly (given suitable processing).

Applications:

- generating fake books, guitar chords
- feeding into models of music cognition, melodic memory

The task: identifying chords and assigning labels in music

- currently to MIDI transcriptions of performances;
- could be applied to audio directly (given suitable processing).

Applications:

- generating fake books, guitar chords
- feeding into models of music cognition, melodic memory

Introduction

Motivation

Modelling

Chords
Structure

Future

Summary

$$p(\mathbf{x}|c; \Omega) = p_D(t\bar{t}|c; \Omega)p_D(rmd|c; \Omega)$$

$$p_D(\mathbf{x}|\alpha_c) = \frac{1}{B(\alpha)} \prod_i x_i^{\alpha_i - 1} \quad \left(\sum_i x_i = 1 \right)$$

$$p(c|\mathbf{x}\Omega) = \frac{p(\mathbf{x}|c\Omega)p(c\Omega)}{\sum_c p(\mathbf{x}|c\Omega)p(c\Omega)}$$

Introduction

Motivation

Modelling

Chords
Structure

Future

Summary

$$p(\mathbf{x}|c; \Omega) = p_D(t\bar{t}|c; \Omega)p_D(rmd|c; \Omega)$$

$$p_D(\mathbf{x}|\alpha_c) = \frac{1}{B(\alpha)} \prod_i x_i^{\alpha_i - 1} \quad \left(\sum_i x_i = 1 \right)$$

$$p(c|\mathbf{x}\Omega) = \frac{p(\mathbf{x}|c\Omega)p(c\Omega)}{\sum_c p(\mathbf{x}|c\Omega)p(c\Omega)}$$

Introduction

Motivation

Modelling

Chords
Structure

Future

Summary

$$p(\mathbf{x}|c; \Omega) = p_D(t\bar{t}|c; \Omega)p_D(rmd|c; \Omega)$$

$$p_D(\mathbf{x}|\alpha_c) = \frac{1}{B(\alpha)} \prod_i x_i^{\alpha_i - 1} \quad \left(\sum_i x_i = 1 \right)$$

$$p(c|\mathbf{x}\Omega) = \frac{p(\mathbf{x}|c\Omega)p(c\Omega)}{\sum_c p(\mathbf{x}|c\Omega)p(c\Omega)}$$

Introduction

Motivation


Modelling

Chords

Structure

Future

Summary



When does one chord end and another begin?

Assumptions:

- *bar* as fundamental unit;
- new chords only on beats.

Our MIDI transcriptions are clear enough

Models: all possible beatwise divisions of a bar. Choose using Bayesian model selection.

$$p(w|\mathbf{x}\Omega') \propto \sum_c p(\mathbf{x}|cw\Omega')p(cw\Omega')$$

Current investigation: how much do experts' opinions on this task differ?

When does one chord end and another begin?

Assumptions:

- *bar* as fundamental unit;
- new chords only on beats.

Our MIDI transcriptions are clear enough

Models: all possible beatwise divisions of a bar. Choose using Bayesian model selection.

$$p(\omega|\mathbf{x}\Omega') \propto \sum_c p(\mathbf{x}|c\omega\Omega')p(c\omega\Omega')$$

Current investigation: how much do experts' opinions on this task differ?

Segmentation of musical audio. Applications:

- Chorus-finding: summary generation
- Structural annotation
- Query-by-content systems

Pop music: timbre / the way it sounds

Simple clustering leads to fragmentation:

- HMM states have exponential duration prior
- one-frame segments

Segmentation of musical audio. Applications:

- Chorus-finding: summary generation
- Structural annotation
- Query-by-content systems

Pop music: timbre / the way it sounds

Simple clustering leads to fragmentation:

- HMM states have exponential duration prior
- one-frame segments

Segmentation of musical audio. Applications:

- Chorus-finding: summary generation
- Structural annotation
- Query-by-content systems


Pop music: timbre / the way it sounds

Simple clustering leads to fragmentation:

- HMM states have exponential duration prior
- one-frame segments


Introduce a duration prior.

$$\varepsilon_{\mathcal{H}}(x, \nu, \gamma) = \frac{1}{|\nu|} x^{-\nu} + (\gamma + 1) \log x.$$



Anneal over candidate segmentations, using modified Wolff algorithm.


Nirvana:Smells Like Teen Spirit, ms(200)/ms(400), 1/12-octave




histseg(wolff,genhyp) : regions(0.0603,0.0424,0.949)




segmentation



annotation



Chumbawamba:Tubthumping, ms(200)/ms(400), 1/12-octave



0 20 40 60 80 100 120 140 160 180

histseg(wolff,genhyp) : regions(0.0903,0.403,0.753)



segmentation



0 20 40 60 80 100 120 140 160 180

time/s

annotation



0 20 40 60 80 100 120 140 160 180

time/s

Investigating recordings

Statistical
Modelling of
Music

Christophe
Rhodes

Introduction


Motivation

Modelling

Chords
Structure

Future

Summary



Musicologists make testable hypotheses, too!

- Tristan Chord (Ernst Knuth)
- Secret Chromatic Art (Edward E. Lowinsky)

Musicologists make testable hypotheses, too!

- Tristan Chord (Ernst Knuth)
- Secret Chromatic Art (Edward E. Lowinsky)

Musicologists make testable hypotheses, too!

- Tristan Chord (Ernst Knuth)
- Secret Chromatic Art (Edward E. Lowinsky)

Music:

- complex activity;
- interactions between creators, performers, listeners, overall culture;
- people play multiple roles.

Statistical modelling:

- building intelligent tools;
- begin to understand the interactions involved