


Where's the chorus? A computational approach for the automatic segmentation of pop songs

Daniel Müllensiefen, David Lewis, Geraint Wiggins
(Goldsmiths College, University of London)


Overview

- The M^{4}S project (*Modeling Music Memory and the Perception of Melodic Similarity*)
- The method: Corpus-based musicology
- The pop music data base
- Lead-sheet computation and segmentation
- Perspectives of automatic analysis


The M^{4}S project: Goals

- Main goal: Address questions about music memory
  - What can ordinary people remember from a tune heard just once?
  - Explain typical memory errors for melodies
  - Determine cognitively relevant features of melodies


The M^{4}S project: Steps

- Automatically analyse and annotate large pop song collection
- Define computable features of pop music (inspired by analytical approaches, e.g. Tagg, Moore, Everett)
- Test features for cognitive validity in psychological experiments
- Describe statistical distributions of features in pop song collection
- Model (implicit) human memory for pop music melodies


The method: Corpus-based musicology

- Determine distribution of musical features in coherent corpus of musical works (patterns of song structure and instrumentation, chord progressions, melodic formulae, rhythm patterns etc.)
- Find statistical associations between features in corpus
- Frequent feature combinations: Musical style prototypes?
- Infrequent feature combinations: Awkward or original cases?


The method: Corpus-based musicology

- Impossible:
  - Detection of subtle differences (find mass not class)
  - Providing (cultural) explanations
  - Detection of aesthetic connotations, meaning and interpretations
- Possible:
  - Tracing musical patterns through pop music history
  - Comparing musical patterns with cultural data (style analysis, chart data, composer background)


The pop music data base

- The collection: 14,067 high-quality transcriptions of full pop songs from 1950s-2006
- Raw data:
  - Polyphonic MIDI files
  - Lyrics
  - Discographical database
  - Chart data


The pop music data base

- Information contained: Compositional structure and arrangement (melodies, harmonies, rhythms, instrumental voices)
- Information only partially contained: Timbre and performance (‘real’ sounds, expressive deviations in rhythm and pitch)

The pop music data base: Example


Lead-sheet computation

- Lead-sheet: Basic representation for comparative analysis of songs
- Components:

a) Identification of all song sections (incl. repetition information)
b) Harmonic progression (chord labeling)
c) Prototypical (monophonic) tune
d) Lyrics


Lead-sheet computation: Song segmentation

- Partition song in meaningful sections, establish relationships between sections and assign functional labels
- Use combination of supervised and unsupervised classification and rule based label assignment in five steps.

Lead-sheet computation: Song segmentation

1. Locate boundaries at points in song where instrumentation changes and that are not too close to each other => List of time points

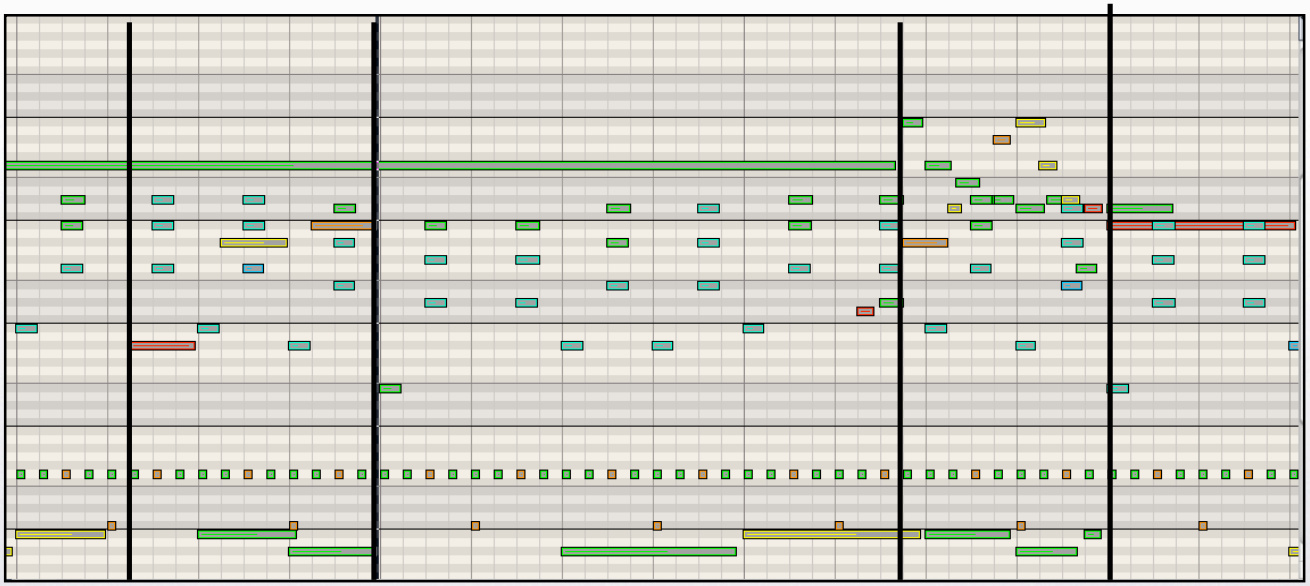

Lead-sheet computation: Song segmentation

2. Find up to five pairs of sections (as demarcated by boundaries) that have same key, same coverage of main melody and are harmonically most similar.

3. Join pairs of sections to clusters if they are harmonically similar enough

Lead-sheet computation: Song segmentation

4. Assign each of the remaining single sections to cluster that is harmonically most similar

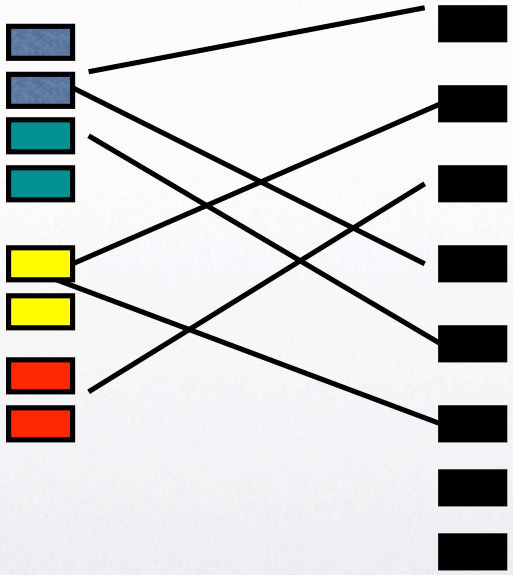

Lead-sheet computation: Song segmentation

5. Label a section cluster

- a chorus if it has mostly repeating lyrics,
- a verse if it has mostly non-repeating lyrics,
- an intro/outro if it appears not more than twice and is starting/ending section, and
- all other clusters else a bridge

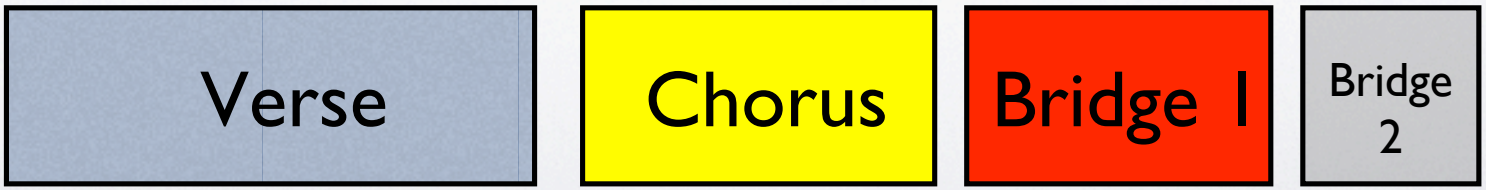


Lead-sheet computation: Song segmentation

- Example: Labeled song sections (at #beats from the beginning) in *Englishman in New York*
- Part 1 Part 2 Part 3 Part 4
- Verse/Intro Verse Verse Chorus
- Beat: 8 40 72 104


Lead-sheet computation: Song segmentation

- Next steps:
  - Include information on chord sequence repetitions ("approximate subsequence matching") for boundary detection
  - Compute complete lead-sheet: Put together information on song structure, harmonic labels, and main tune identification
  - Identify famous structure prototypes, e.g.
    - Blues / Ballads: A A A A
    - Variants of Verse-Chorus: A B A B (C) A B B B
    - Rhapsodian form: A B C D E F
    - Vamp based forms: A___ B A_____ B
  - Classify frequent but unnamed forms (= learn from data)

Perspectives of automatic pop song analysis

- A) Find generic patterns / formulae, e.g.: Most frequent harmonic patterns
- 52% of all chord sequences used in database are combinations of **I IV** and **V** chords and
- Very frequent turnaround sequence (see Moore, 2006; Kramarz, 2007: **I vi IV V**) in 5% of sequences, and in 16% of all 14000 songs
- Most frequent after **I IV V** combinations are alternations between **I** and **vi** and **I** and **ii** (each 3.5% of chord sequences)
- Specific sequences are much less frequent …

Perspectives of automatic pop song analysis

- A) Find generic patterns / formulae, e.g.: Most frequent melodic patterns
  - Most frequent melodic phrases, pitch intervals only:
  - Note repetitions (#1)

  - Combinations of note repetitions and seconds (#5)

  - Descending major scale: first pattern with range >maj3 (#35)

Perspectives of automatic pop song analysis

- B) Find particular patterns (motifs/riffs) through music history e.g.:
  - Melodic patterns, e.g. opening phrase of *Brown girl in the ring*:

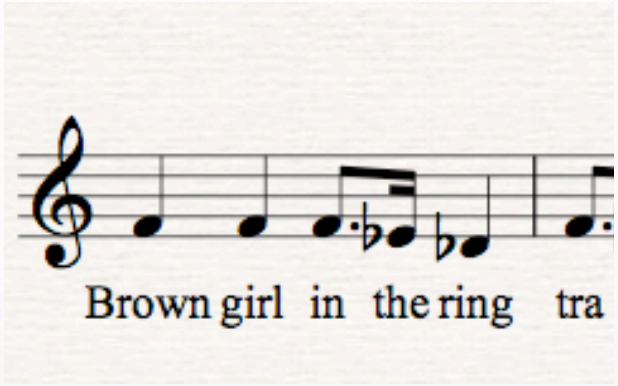

  - Represented as tuples of intervals and duratio ratios: (0,1) (0,.75) (-2,.3) (-2,4) (+4,.75)
  - => Occurs in 26 different songs in database, e.g. Rick Astley: *Together forever*; The Beatles: *Maxwell's silver hammer*; Bon Jovi: *In these arms*
  - => Any relations to Jamaican children's rhyme or Boney M. song?

Perspectives of automatic pop song analysis

- B) Find particular patterns through music history e.g.: Harmonic patterns, e.g. opening chord sequence from *Yesterday:* **I vii III vi IV**
  - only in 14 songs, e.g. Make me smile (Chicago), Sara (Jefferson Starship), Yesterday (Wet Wet Wet)
  - only in one song by The Beatles
  - relatively common in jazz standards
  - Similarly: Find rhythmical patterns, e.g. the Bo Diddley riff: || **1**+2**+**3+**4**+ | 1+**2**+**3**+4+ :||
- <u>Potential of approach:</u>
  - Give quantified structural descriptions of pop music styles and identify style prototypes
  - Write history of composition in pop music through re-use of patterns and pattern associations (co-occurences)


Next steps

- 1. Use more sophisticated algorithms for
- the generation of descriptive features, e.g. melody contours (Frieler et al., 2007), melodic accents (Pfleiderer et al., 2007), rhythm patterns, complexity (Pearce and Wiggins, 2004)
- and for
- the establishment of similarity relations between musical entities, e.g. melodic, harmonic, and rhythmic similarity (Müllensiefen and Frieler, 2004, 2007)
- 2. Test cognitive validity of algorithmic features
- 3. Address more interesting questions from a pop music analysis perspective


Questions and invitation for collaboration

- Good structural definition of song segments?
- Classification of song structures?
- Ideas for summarising chord sequences?
- What are interesting patterns to look for in the collection? (Melodic motifs, rhythm patterns, chord progressions)
- Connections to existing analytical research?


Where's the chorus? A computational approach for the automatic segmentation of pop songs

Daniel Müllensiefen, David Lewis, Geraint Wiggins
(Goldsmiths College, University of London)