School of informatics

Music Informatics

Alan Smaill

Mar 25 2014

Today

School of informatics

- Lerdahl and Jackendoff, GTTM grammar modules
- HighC, quick reminder
- Examinable material

Recall GTTM

Recall Lerdahl and Jackendoff's "Generative Theory of Tonal Music" (GTTM) –
an attempt to characterise tonal music via a characterisation of how a listener familiar with the style makes "musical sense" of music in these styles.

We saw before the Grouping level, which provides ideas for chunking material. Today look again at this proposal, and the other levels described in GTTM.
Handout is from Jackendoff's "Consciousness and the Computational Mind", chapter 11 "Levels of Musical Structure" and gives a good overview.

Resources

School of informatics

While unfortunately the handout chapter is not in the public domain on-line, two valuable works are:

- GTTM itself (there is a lot of detail)

  `http://cognet.mit.edu/library/books/view?isbn=026262107X`

  and

- Chapter "musical parsing and musical affect" from "Languages of the Mind":

  `http://cognet.mit.edu/library/books/chapter?isbn=0262100479&part=chap7`

Musical Surface

School of informatics

- Reminder of choice of musical surface as (roughly) notes in GTTM;
- More generally:

> *Hence a full psychological theory of music must account for the derivation of the musical surface from an acoustic signal. the musical surface, however, is the lowest level of representation that has musical significance.*
>
> *Jackendoff, Consciousness and the Computational Mind, p 219*

Levels of Musical Structure

School of informatics

These are built on the surface level, which is itself sequential, in hierarchical fashion.
The lowest level is that of **Grouping Structure**, which we already saw, with associated rules of where boundaries may occur. The Mozart example used has a history going back to Leonard Bernstein's proposals in "The Unanswered Question: Six Talks at Harvard", 1976.

The grouping analysis is itself hierarchical — groups of notes, then groups of groups, and so on. The rules as given allow many possible parses.

The examples of "good" and "bad" parses are meant to strengthen the claim that the preference rules in GTTM do point the way to the musically significant analyses.

Metrical Structure

We have seen approaches to metrical structure before, and GTTM at this level is similar to earlier approaches. Again, the metrical level is itself hierarchical. The handout mentions empirical evaluation of GTTM's specific rules for metrical structure.

Note that the metrical structure is defined only up to a relatively small duration (eg two bars) — this is unlike grouping structure, which can very well scale up to larger and larger groups.

Metrical example

The metrical level gives us an analysis like the following. Here, grouping is up to the level of two bars — but note that the metrically stronger of the bars corresponds to the start of the **second** full bar. (In fact this is the third bar in the score, because there are 3 beats of background "vamp" before the melody enters).

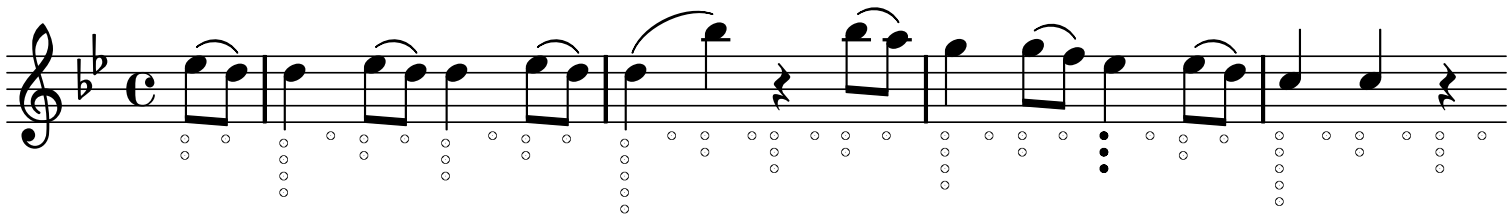

This **can** be heard or played with the opposite two-bar phase – but L&J (and Bernstein before) argue this is the right version.

Corresponding grouping analysis

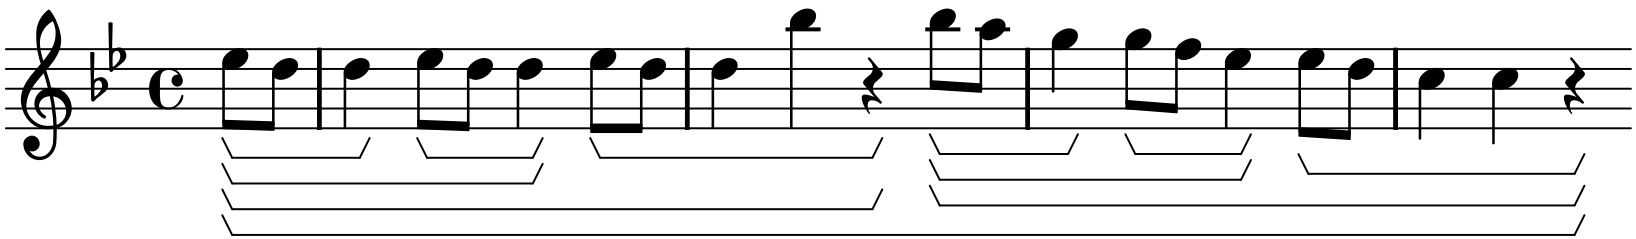

Similar, but not identical . . .

Between grouping and metrical levels

School of informatics

Compare the grouping and the metrical parses:

- There is agreement on the significance of the two bar length.
- The grouping analysis recognises the four bar grouping.
- The boundaries of the two bar length chunks are not precisely aligned (they are out of phase).

Thus we end up with related, but distinct, decompositions of the musical surface. L&J claim these are both cognitively significant. We can see that the parsing task is getting complicated!

Correspondence Rules

L&J use some **correspondence rules** to cover the relationship between different levels of the analysis. These include rules already seen, like the grouping **symmetry** rule:

> *. . . prefer groupings that respect musical parallels . . .*

where "parallel" may refer to metrical similarity.

More explicitly, L&J introduce a grouping rule referring to correspondence between the grouping level and the two levels we have yet to see:

> **GPR 7 (Time-Span and Prolongational Stability)**
> *Prefer a grouping structure that results in more stable time-span and/or prolongational reductions.*

The first two levels address primarily the rhythmic aspect of music. The **Time-span** level looks at the pitch and harmonic information, so as to regard some passages as a form of **elaboration** of others; an example is where a simple melody is decorated, or variations on it are built.

The claim here is the following:

> **Reduction Hypothesis**
> *The pitch-events of a piece are heard in a hierarchy of relative importance; structurally less important events are heard as ornamentations or elaborations of events of greater importance.*

Example

School of informatics

See the example in the handout;
the parse tree at the top indicates the analysis in terms of which groups are considered as taking priority at different levels of the reduction.

So, in a theme with variations, or considering whether aspects like ornamentation change whether or not we are listening to the same piece, look for a common structural underpinning in the form of the time-span reduction.

Example with three levels

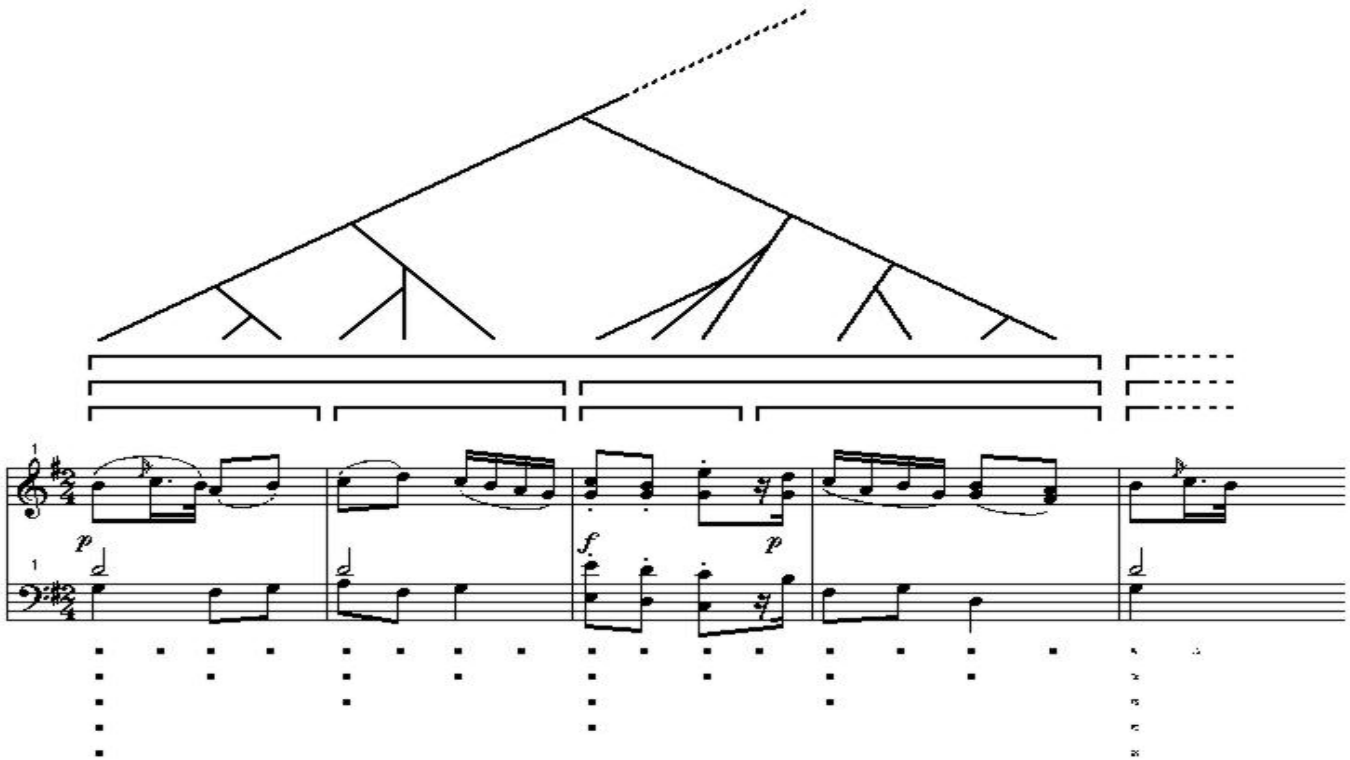

Figure 1(a)

Image due to Ian Cross, in Music Analysis, 1998, vol 17. No.1

`www.jstor.org/stable/854368`

Time span rules

School of informatics

To get an idea of how the time-span tree is formed, look at sample rule from GTTM.
The **head** of a reduction is the part which is considered the more fundamental – others are subordinate.
Harmonic preference depends on a notion of which harmonies are more or less consonant, and how closely they relate to the tonic in a given key.

> **TSPR 2 (Local Harmony)** *Of the possible choices for head of a time-span T, prefer a choice that is*
>
> 1. *relatively intrinsically consonant,*
> 2. *relatively closely related to the local tonic.*

Time-span ctd

School of informatics

Notice that this makes the time-span reduction closely dependent on the classical tonal aspects of the style addressed (unlike grouping, and to some extent the metrical level). Any implementation has to address the whole language of key, cadence and so on.

Implementations of the GTTM grammar have largely focused on the first two levels, in practice.

Prolongational Reduction

School of informatics

This level builds on the previous levels, and is intended to capture the notion of increase or decrease of **tension** — in the passage from one analysed component to another, is the music heard as **confirming** or **denying** the material in the first component?
See handout discussion for this also.

Notice the dependencies in the processing implicit in the diagram of how the levels are detected.
In particular, there is coupling between grouping and metrical levels, and between time-span reduction and prolongational reduction.

HighC

Reminder of the graphical compositional tools UPIC and HighC now that we have seen Xenakis's diagrammatic layout which preceded deployment as conventional scores.

In both cases (though HighC provides more functionality):

- Main drawing area treats time left-to-right, pitch bottom to top;
- Pitched sounds are drawn on the "canvas";
- Different colours correspond to different timbres;
- A subsidiary area allows sounds to be designed themselves before use;
- ... and to draw volume (intensity) envelopes for the given section.

view of UPIC

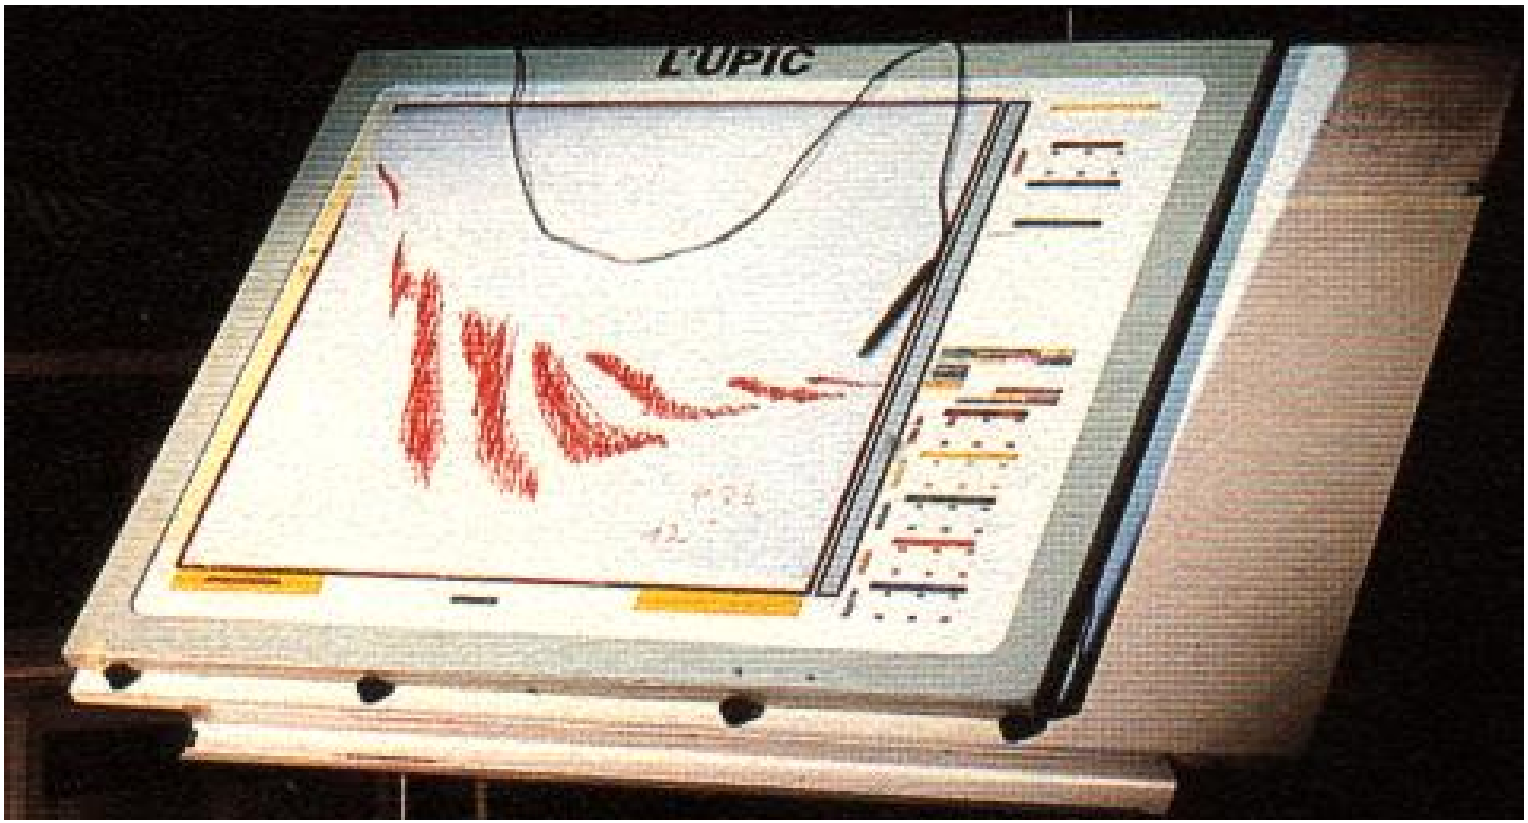

View of HighC

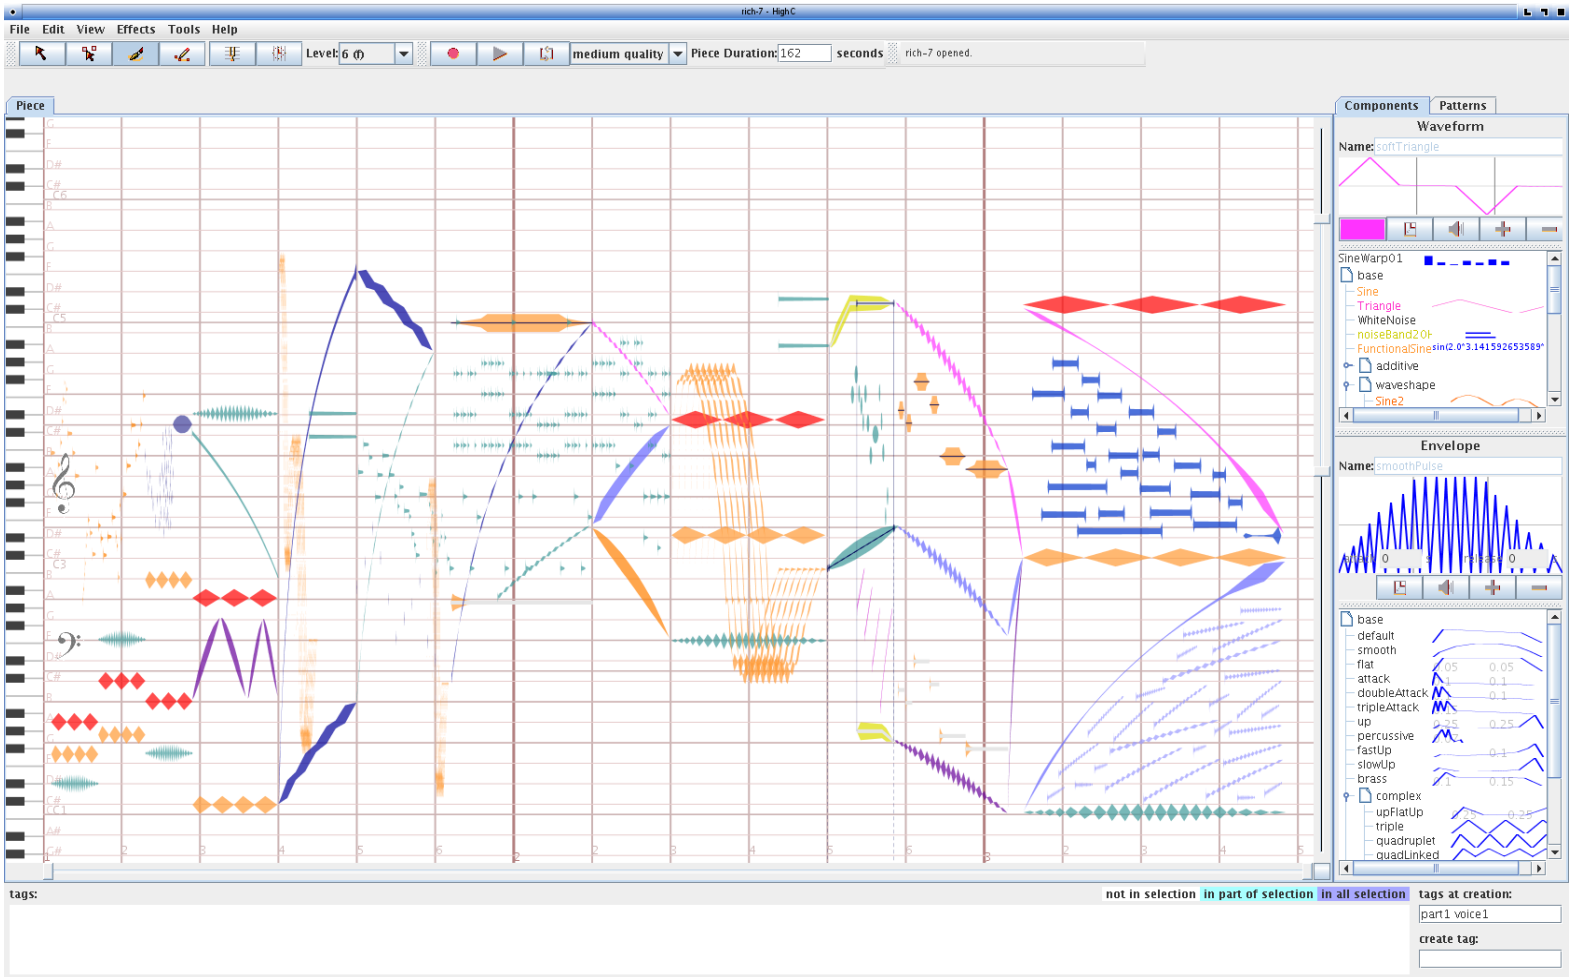

Examinable material

School of informatics

- Basics of different representations (midi, WTM score, mp3, sound-file):
  strengths and weaknesses, what they make easy and what they make hard
- Converting between representations (midi <-> mp3/wav), what's involved (without technical details)
- WTM metrical hierarchy, what it is.
- Longuet-Higgins metrical analysis algorithm
  main idea, know some of the rules (not all)
- musical ambiguity, how recognise? how deal with by machine?
- 2-d pitch array, how use to recognise key and note spelling

Examinable material ctd

School of informatics

- Beat-tracking: what the problem is, Dixon's algorithm (in outline)
- Score-following: what the problem is; outline HMM approach; the main architecture of Raphael's system
- Musical grammars, what they look like, what they can be used for
- GTTM: what it tried to do, the main components, cognitive claim
- GTTM grouping rules: some examples of the rules; what choices need to be made to implement these rules.
- Rule-based systems: what they are, what needs to be put in, how control the use of the rules

Examinable material ctd

School of informatics

- paradigmatic analysis: what it is, algorithm for carrying it out (parametrised on notion of similarity)
- Cope's style imitation: how he did it, notion of signature
- Multiple viewpoints: several viewpoints, what difference it makes, what entropy calculation tells us
- Pitch-class sets: what they are, when two sets of pitches are the same pitch-class. Why the theory was developed.

Examinable material ctd

School of informatics

- Xenakis: how generate music using statistical ideas, how generate musical curves from straight lines.
  UPIC and successor system – human interface to the system.
- Components of improvisation system in OMax
- Learning at different levels; what is a probabilistic grammar

Summary

School of informatics

THE UNIVERSITY OF EDINBURGH

- GTTM: the four levels.
- Brief reminder of graphical interface to composition tools
- Examinable material