


Mining for the Meaning of Music

Dan Ellis

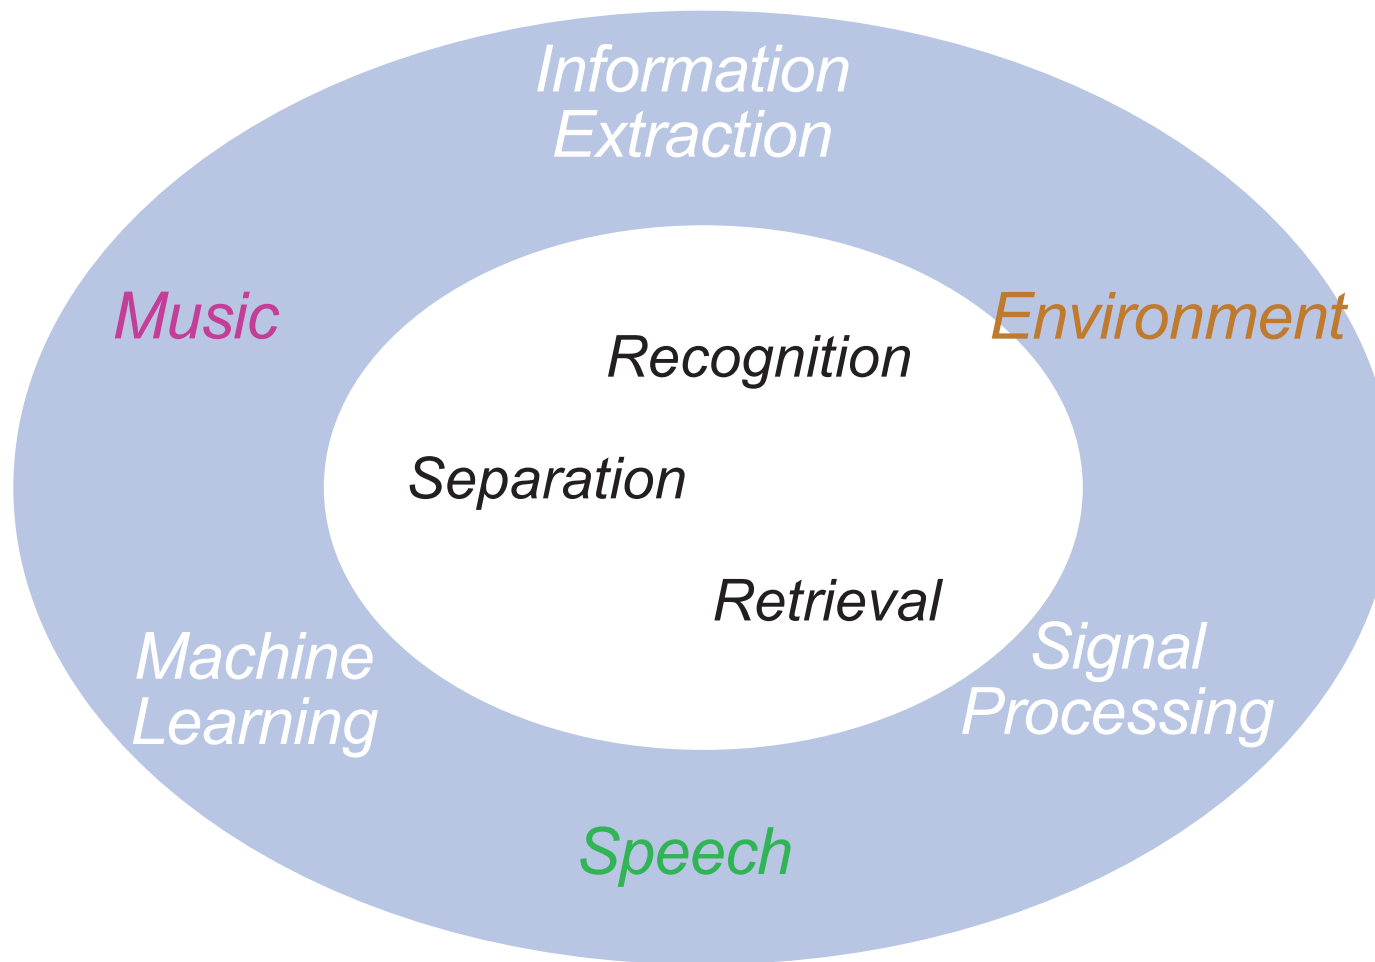
Laboratory for Recognition and Organization of Speech and Audio
Dept. Electrical Engineering, Columbia University, NY USA

<http://labrosa.ee.columbia.edu/>

1. Motivation: What is Music?
2. Eigenrhythms
3. Melodic-Harmonic Fragments
4. Example Applications



LabROSA Overview



I. Motivation: What is music?

- What does music **evoke** in a listener's **mind**?
- Which are the **things** that we call “**music**”?



Oodles of Music





- What can you do with a million tracks?

Re-use in Music


- What are the most **popular** **chord progressions** in pop music?

Scatter of PCA(3:6) of 12x16 beatchroma



Potential Applications


- Compression
- Judgments / classification
- Manipulation



2. Eigenrhythms: Drum Track Structure

Ellis & Arroyo ISMIR'04

- To first order,
all pop music has the **same beat**:




- Can we **learn** this from **examples**?

Basis Sets


- Combine a few basic patterns to make a larger dataset

$$\text{data } X = \text{weights } W \times \text{patterns } H$$




Different basis projections


Principal Component Analysis (PCA)




Independent Component Analysis (ICA)




Linear Discriminant Analysis (LDA)



Nonnegative Matrix Factorization (NMF)




Drum Pattern Data



- Tempo normalization + downbeat **alignment**

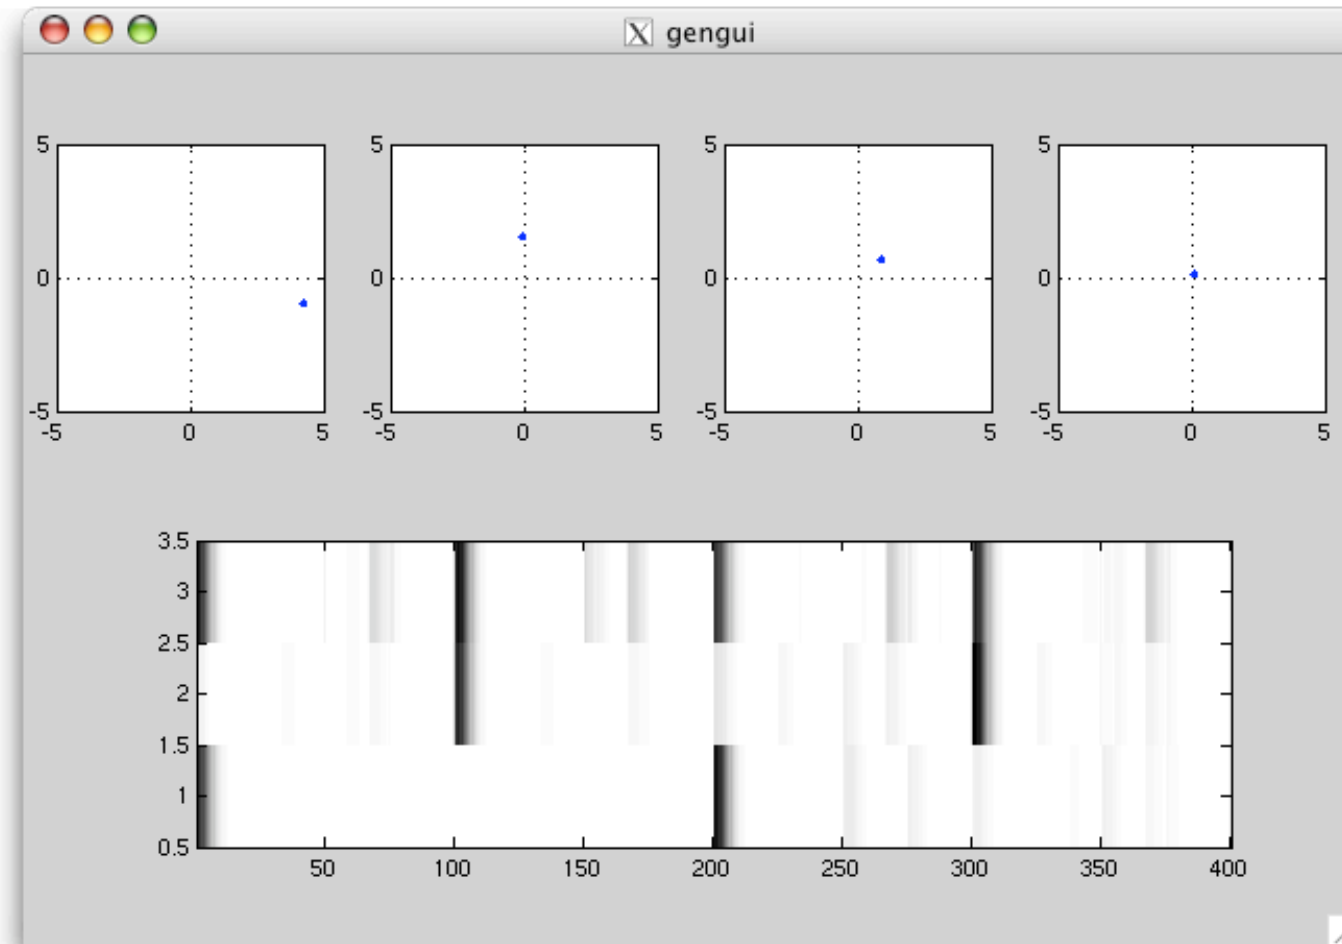
NMF Eigenrhythms



- Nonnegative: only **add** beat-weight


Eigenrhythm BeatBox

- Resynthesize rhythms from eigen-space




3. Melodic-Harmonic Fragments

- How **similar** are two pieces?




- Can we find all the pop-music **clichés**?
 - | - V - VI - IV




MFCC Features

- Used in **speech recognition**


Let It Be - log-freq specgram (LIB-1)



MFCCs




Noise excited MFCC resynthesis (LIB-2)




Chroma Features

- To capture “musical” content


Let It Be - log-freq specgram (LIB-1)




Chroma features



Shepard tone resynthesis of chroma (LIB-3)




MFCC-filtered shepard tones (LIB-4)




Beat-Synchronous Chroma


Let It Be - log-freq specgram (LIB-1)




Onset envelope + beat times



Beat-synchronous chroma



Beat-synchronous chroma + Shepard resynthesis (LIB-6)




Finding Cover Songs


Ellis & Poliner '07

- Little similarity in surface audio...



Let It Be - The Beatles



Let It Be - Nick Cave




- .. but appears in beat-chroma



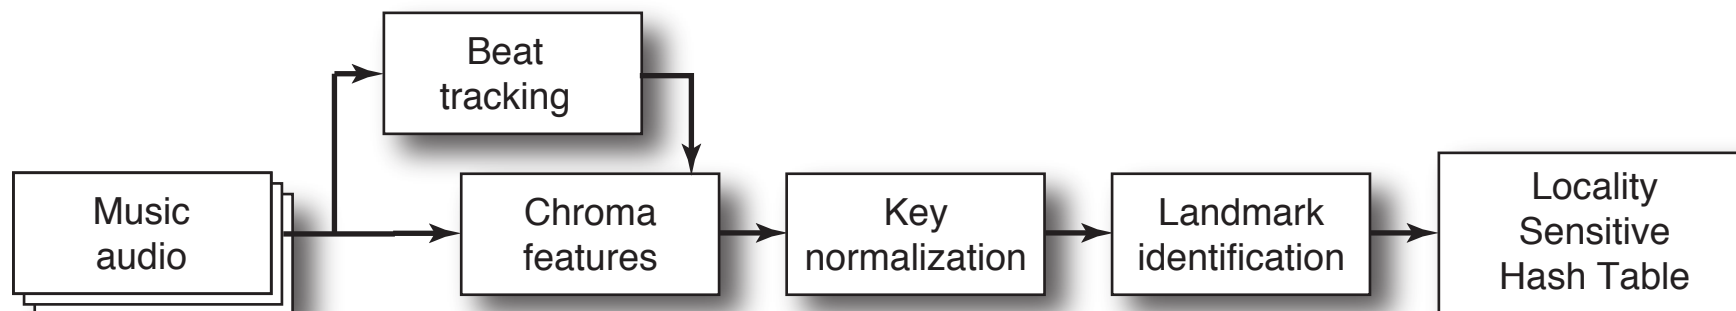
Finding Cover Songs

- Match via global cross-correlation




Finding Common Fragments

- Data mining in beat-chroma database




- .. a large set of excerpts (patches)



Landmarks


- Extract patches from starts of phrases

Come Together - Spectrogram, Beat-sync chromogram, and top 10 segment points



Locality Sensitive Hash


- Nearby items go into same hash bucket via multiple random projections



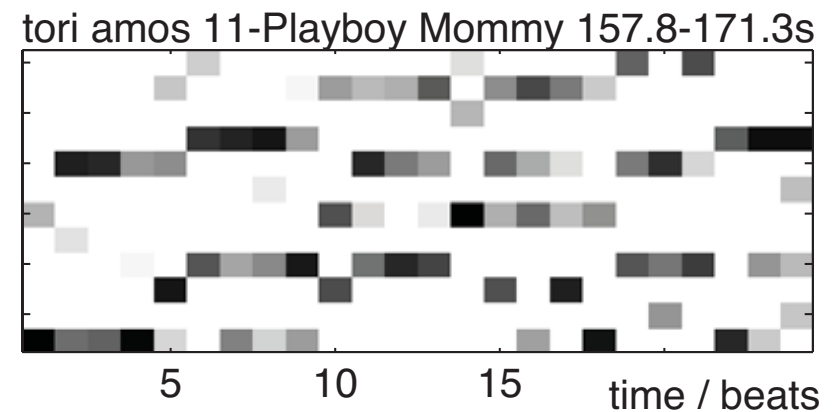



from Slaney & Casey '08

Common Fragments

- **Cluster** beat-synchronous chroma patches




Clustered Fragments



- ... for a **dictionary** of common themes?


NMF Decomposition

Beat-chroma fragments: 30 Basis NMF decomposition



NMF Decomposition

- Originals and 30-basis NMF reconstructions



- Conclusion?

4. Example Applications: Music Discovery

Berenzweig & Ellis '03

- Connecting listeners to musicians

Playola Search: Artist
[\[About\]](#) [\[Help\]](#) [\[Turn Samples Off\]](#) [\[Turn Debug On\]](#) [\[Turn Poupus Off\]](#) [\[Logout dpwe\]](#)

Get Playola Selections: 20 songs you recently heard Browse: [Artists](#) [Albums](#) [Playlists](#) Range: 0-C

Artist: [The Woodbury Muffin Outbreak](#) [\[band web page\]](#) [\[Play!\]](#) Playlist: -New Playlist- [\[Add to\]](#) [\[View\]](#)

	Song Title	Artist	Time	Rating
<input type="checkbox"/>	The Ballad of Tabitha	The Woodbury Muffin Outbreak	4:00	<input type="checkbox"/>
<input type="checkbox"/>	Monkey Dreams	The Woodbury Muffin Outbreak	2:57	<input type="checkbox"/>
<input type="checkbox"/>	A Cold Dark Night (Live)	The Woodbury Muffin Outbreak	3:13	<input type="checkbox"/>
<input type="checkbox"/>	Leo, The Ballad of	The Woodbury Muffin Outbreak	1:48	<input type="checkbox"/>
<input type="checkbox"/>	Baby I Forgot To Tell You	The Woodbury Muffin Outbreak	4:04	<input type="checkbox"/>

Music-Space Browser [\[What's This?\]](#)

Feature	Less	More
AltNGrunge	<input type="checkbox"/>	<input type="checkbox"/>
CollegeRock	<input type="checkbox"/>	<input type="checkbox"/>
Country	<input type="checkbox"/>	<input type="checkbox"/>
DanceRock	<input type="checkbox"/>	<input type="checkbox"/>
Electronica	<input type="checkbox"/>	<input type="checkbox"/>
MetalNPunk	<input type="checkbox"/>	<input type="checkbox"/>
NewWave	<input type="checkbox"/>	<input type="checkbox"/>
Rap	<input type="checkbox"/>	<input type="checkbox"/>
RnBSoul	<input type="checkbox"/>	<input type="checkbox"/>
SingerSongwriter	<input type="checkbox"/>	<input type="checkbox"/>
SoftRock	<input type="checkbox"/>	<input type="checkbox"/>
TradRock	<input type="checkbox"/>	<input type="checkbox"/>
Female	<input type="checkbox"/>	<input type="checkbox"/>
HiFi	<input type="checkbox"/>	<input type="checkbox"/>

Similar Songs: [\[Play this list\]](#) [\[What's This?\]](#)

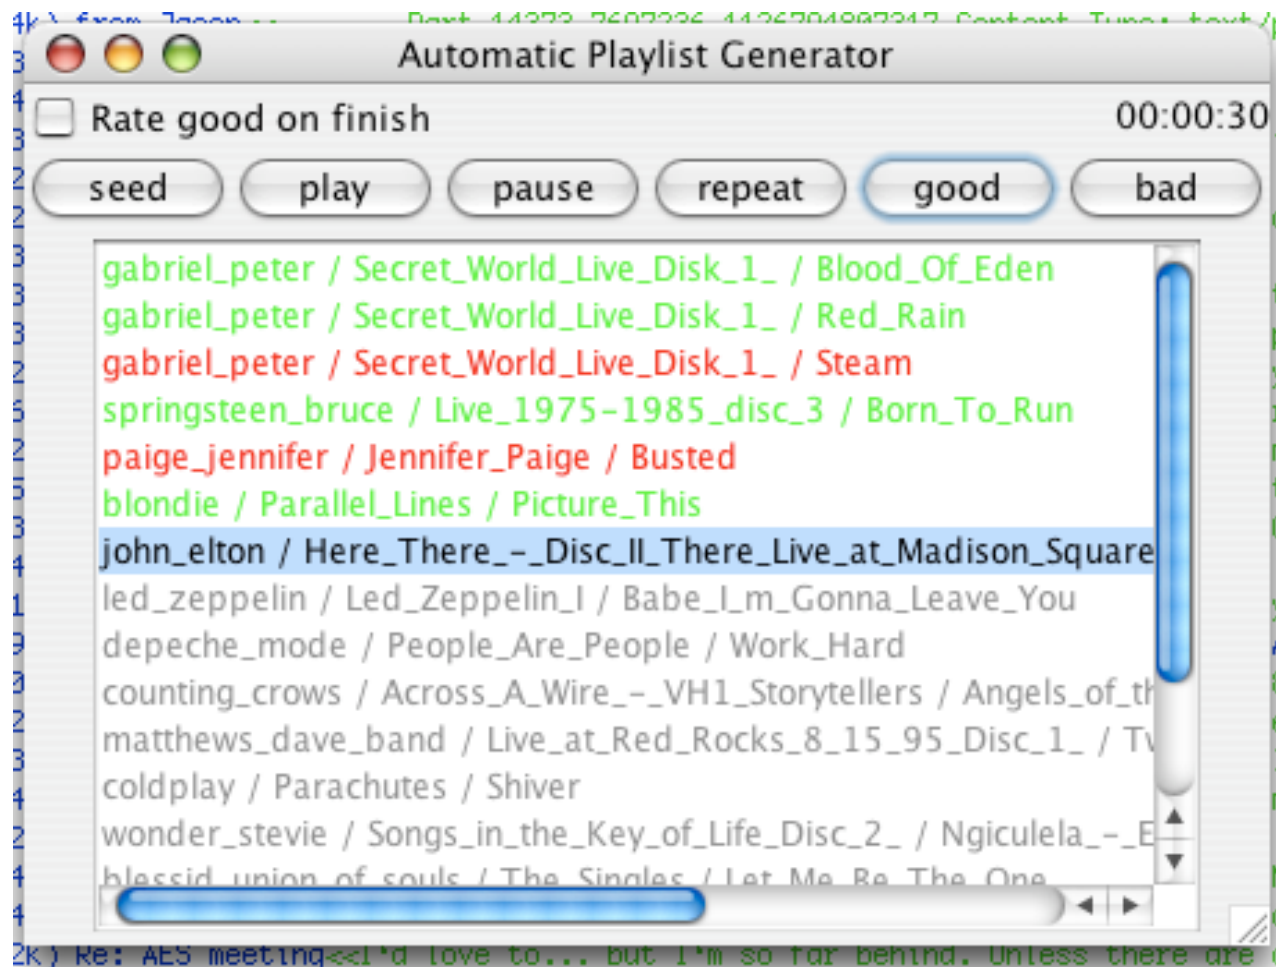
	Song Title	Artist	Distance	Good Match?
<input type="checkbox"/>	Baby I Forgot To Tell You	The Woodbury Muffin Outbreak	0.00	<input type="checkbox"/>
<input type="checkbox"/>	Number five	Bizi Chyld	0.07	<input type="checkbox"/>
<input type="checkbox"/>	Waiting for Your Love	Toto	0.08	<input type="checkbox"/>
<input type="checkbox"/>	Excerpt from 'CD'	Weirdomusic	0.08	<input type="checkbox"/>



Playlist Generation

Mandel, Poliner, Ellis '06

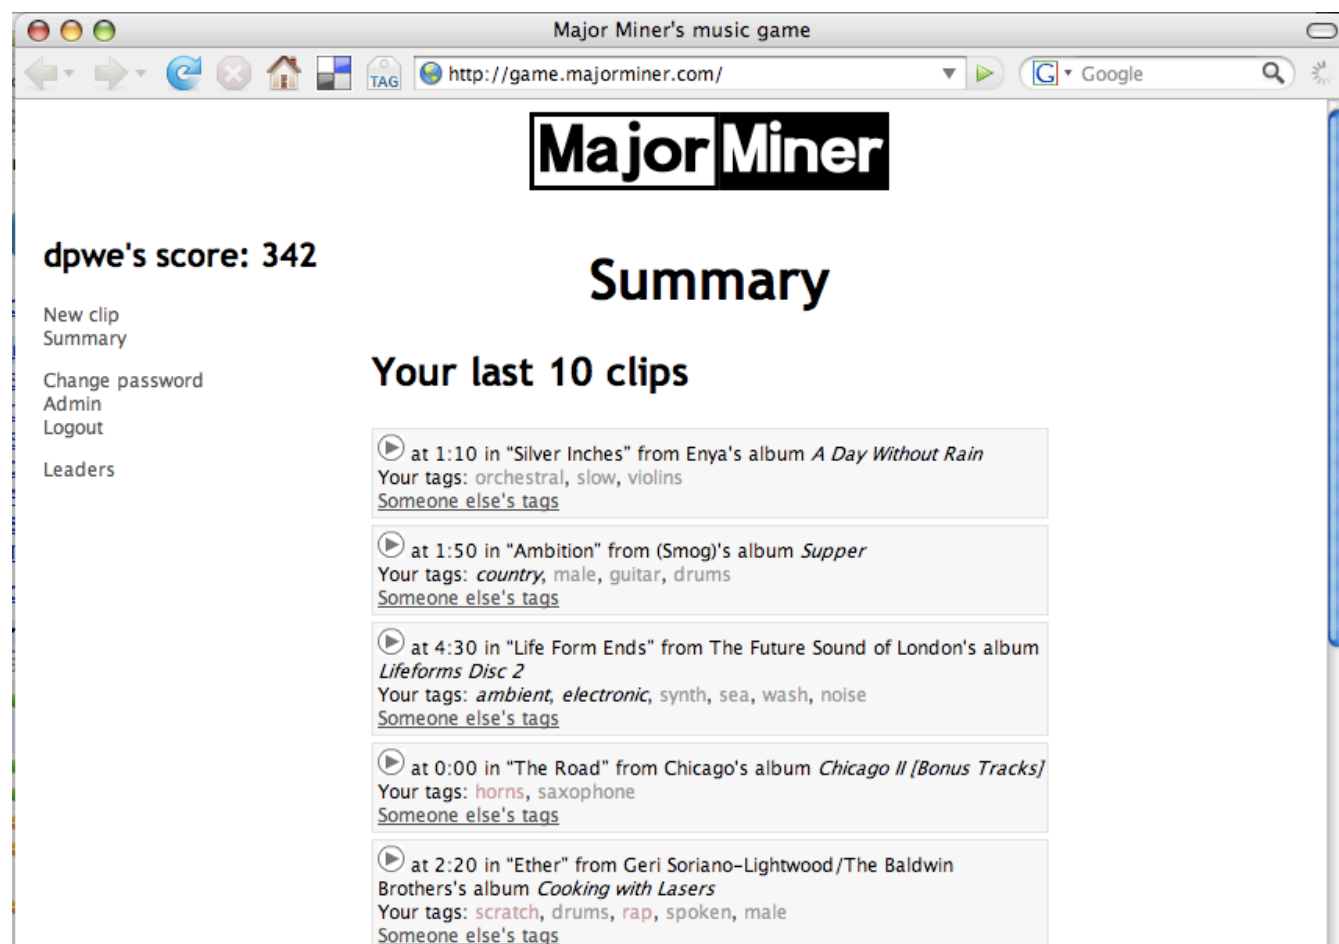
- Incremental learning of listeners' preferences



MajorMiner: Music Tagging


Mandel & Ellis '07,'08

- Describe music using **words**



The screenshot shows a web browser window titled "Major Miner's music game" with the URL "http://game.majorminer.com/". The page features the "Major Miner" logo at the top. On the left, a sidebar lists navigation options: "New clip", "Summary", "Change password", "Admin", "Logout", and "Leaders". The main content area displays "dpwe's score: 342" and a "Summary" section titled "Your last 10 clips". This section lists five clips with their respective tags and a link to "Someone else's tags".

Clip Title	Album	Your tags	Someone else's tags
at 1:10 in "Silver Inches"	Enya's album <i>A Day Without Rain</i>	orchestral, slow, violins	Someone else's tags
at 1:50 in "Ambition"	(Smog)'s album <i>Supper</i>	country, male, guitar, drums	Someone else's tags
at 4:30 in "Life Form Ends"	The Future Sound of London's album <i>Lifeforms Disc 2</i>	ambient, electronic, synth, sea, wash, noise	Someone else's tags
at 0:00 in "The Road"	Chicago's album <i>Chicago II [Bonus Tracks]</i>	horns, saxophone	Someone else's tags
at 2:20 in "Ether"	Geri Soriano-Lightwood/The Baldwin Brothers's album <i>Cooking with Lasers</i>	scratch, drums, rap, spoken, male	Someone else's tags



Music Transcription

Poliner & Ellis '05,'06,'07

Training data and features:

- MIDI, multi-track recordings, playback piano, & resampled audio (less than 28 mins of train audio).
- Normalized magnitude STFT.



Classification:


- N-binary SVMs (one for ea. note).
- Independent frame-level classification on 10 ms grid.
- Dist. to class bndy as posterior.




Temporal Smoothing:

- Two state (on/off) independent HMM for ea. note. Parameters learned from training data.
- Find Viterbi sequence for ea. note.


feature representation



classification posteriors




hmm smoothing



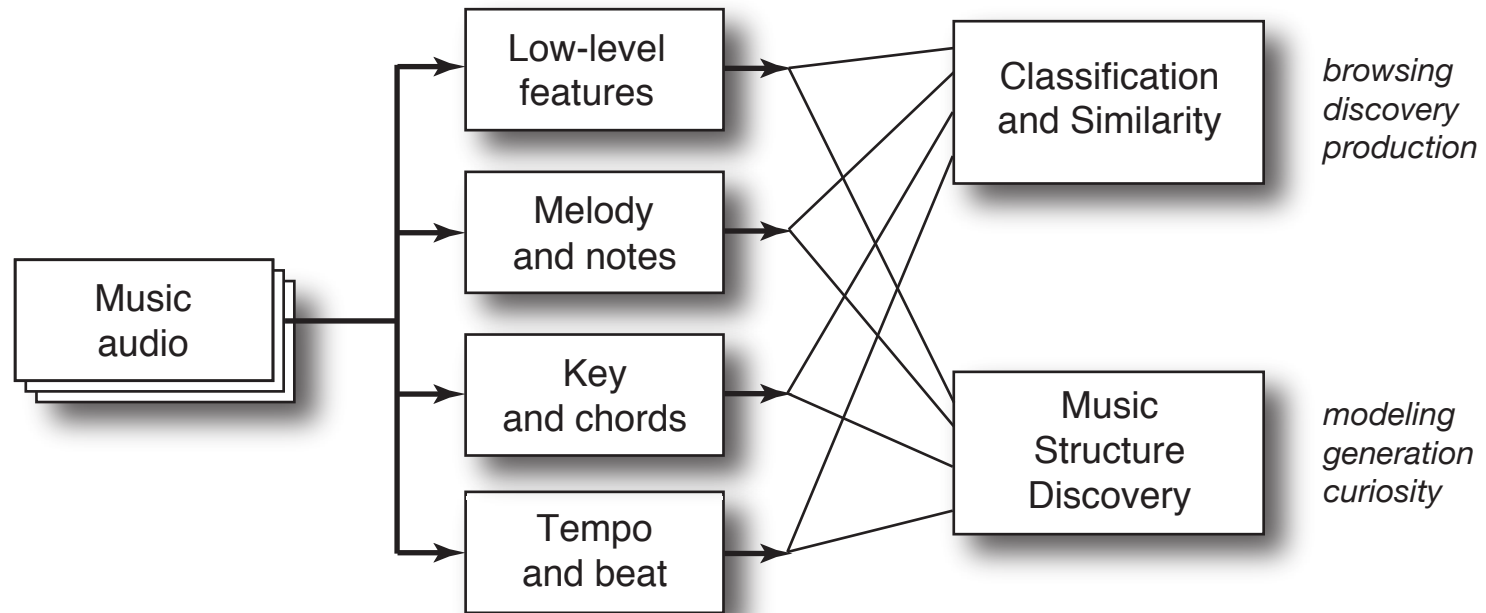
MEAPsoft

- **M**usic **E**ngineering **A**rt **P**rojects
 - collaboration between EE and Computer Music Center

*with Douglas Repetto,
Ron Weiss, and the rest
of the MEAP team*



Conclusions



- Lots of **data**
+ noisy **transcription**
+ weak **clustering**
⇒ musical **insights?**